RUGGEDCOM®
INDUSTRIAL STRENGTH NETWORKS™

ROX®

Rugged Operating System™
on Linux

v2.2 Web Interface User Guide

For RuggedBackbone™ RX1500

February 22, 2012

ROX™: Web Interface User Guide

Copyright © 2011 RuggedCom Inc.

ALL RIGHTS RESERVED

Dissemination or reproduction of this document, or evaluation and communication of its contents, is not authorized except where expressly permitted. Violations are liable for damages. All rights reserved, particularly for the purposes of patent application or trademark registration.

This document contains proprietary information, which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of RuggedCom Inc.

Disclaimer Of Liability

We have checked the contents of this manual against the hardware and software described. However, deviations from the description cannot be completely ruled out.

RuggedCom shall not be liable for any errors or omissions contained herein or for consequential damages in connection with the furnishing, performance, or use of this material.

The information given in this document is reviewed regularly and any necessary corrections will be included in subsequent editions. We appreciate any suggested improvements. We reserve the right to make technical improvements without notice.

Registered Trademarks

RuggedServer™, RuggedWireless™, RuggedCom Discovery Protocol™ (RCDP™), RuggedExplorer™, Enhanced Rapid Spanning Tree Protocol™ (eRSTP™), ROX™, Rugged Operating System On Linux™, RuggedBackbone™ are trademarks of RuggedCom Inc. Rugged Operating System® (ROS®) and RuggedSwitch® are registered trademarks of RuggedCom Inc. Other designations in this manual might be trademarks whose use by third parties for their own purposes would infringe the rights of the owner.

Warranty

Five (5) years from date of purchase, return to factory. For warranty details, visit www.ruggedcom.com or contact your customer service representative.

Contacting RuggedCom

Corporate Headquarters	US Headquarters	Europe Headquarters
RuggedCom Inc. 300 Applewood Crescent Concord, Ontario Canada, L4K 5C7 Tel: +1 905 856 5288 Fax: +1 905 856 1995 Toll-free: 1 888 264 0006	RuggedCom 1930 Harrison St., Suite 209 Hollywood, Florida USA, 33020 Tel: +1 954 922 7938 ext. 103 Fax: +1 954 922 7984 Toll-free: 1 888 264 0006	RuggedCom Unit 41, Aztec Centre, Aztec West, Almondsbury, Bristol United Kingdom BS32 4TD Tel: +44 1454 203 404 Fax: +44 1454 203 403
Email: RuggedSales@RuggedCom.com		

Technical Support
Toll Free (North America): 1 866 922 7975 International: +1 905 856 5288 Email: Support@RuggedCom.com

Web: www.RuggedCom.com

Table of Contents

Preface 25
Supported Platforms 25
Who Should Use This User Guide 25
How This Guide Is Organized 25
Applicable Operating System Software Revision 25
I. Administration 26
1. The ROX™ Web Interface 27
1.1. Getting Started 27
1.1.1. Requirements 27
1.1.2. Connecting To The Web Interface 27
1.1.3. The Web Browser Connection 27
1.2. The Structure Of The Web Interface 28
1.2.1. Top-level Menu Categories 30
1.3. Making Configuration Changes 31
1.3.1. Configuring Tables Using Key Settings Forms 33
1.3.2. Viewing More Information in Tables 35
2. System Administration 37
2.1. Administration menu 37
2.2. System Commands 37
2.3. Administrative Access Control 41
2.4. User Accounts 45
2.5. Software Upgrade 47
2.6. ROXflash Cross-Partition Imaging Tool - Software Downgrade 50
2.6.1. Uses 50
2.6.2. ROXflash Configuration 50
2.7. Scheduling Jobs 52
2.8. The Featurekey 55
2.8.1. Overview 55
2.8.2. Upgrading Feature Levels in the field 55
2.8.3. When a File-based featurekey does not Match the Hardware 55
2.8.4. Viewing RuggedCom Serial Numbers 56
2.8.5. Uploading a Featurekey 57
2.8.6. Backing Up a Featurekey Using the Web User Interface 58
2.9. Installing and Backing Up Files 59
2.9.1. Installing Files 59
2.9.2. Backing Up Files 60
2.10. Deleting Log Files 61
2.11. Saving Full Configurations 61
2.12. Loading Full Configurations 62
3. Time Synchronization 63
3.1. NTP Fundamentals 63
3.1.1. The NTP Sanity Limit 63
3.2. Configuring Time Synchronization 64
3.2.1. Configuring the System Time and Date 64
3.2.2. Configuring the System Time Zone 64
3.2.3. Configuring the Local Time Settings 65
3.2.4. Configuring NTP Servers 65
3.2.5. Adding Server Keys 67
3.2.6. Configuring NTP Server Restrictions 67
3.2.7. Configuring an NTP Server using Multicast or Broadcast 69
3.2.8. Configuring an NTP Client using Multicast 70

3.2.9. Configuring an NTP Client using Broadcast 70
3.2.10. Checking NTP Status 71
4. Basic Network Configuration 72
4.1. IP Interfaces 72
4.1.1. Configuring an IP Address 72
4.1.2. Simple Network Setup with the Default IPv4 Addresses 73
4.1.3. Configuring an IPv6 Address 74
4.1.4. Simple Network Setup with IPv6 Addresses 75
4.1.5. Routable Interfaces 76
5. IP Network Interfaces 77
5.1. IPv6 Fundamentals 77
5.1.1. Addressing 77
5.1.2. Security 77
5.1.3. IPv6 Address Scopes 77
5.1.4. IPv6 Multicast Addresses 77
5.2. IPv6 Neighbor Discovery 78
5.3. Adding Interfaces to Switched Ports 81
5.3.1. All-VLANs 83
5.4. Non-switched Interface Menu 85
5.4.1. Configuring IP Address Source and ProxyARP for Non-switched Interfaces 87
6. Alarms 89
6.1. Introduction 89
6.1.1. Alarm Subsystems 89
6.1.2. Fail-Relay Behavior 89
6.1.3. Alarm LED Behavior 89
6.1.4. Clearing and Acknowledging Alarms 89
6.2. Alarm Configuration 90
6.2.1. Administrative Alarm Configuration 93
6.2.2. Chassis Alarm Configuration 94
6.2.3. Switch Alarm Configuration 95
7. Domain Name Search 96
7.1. Domain Name Lookup 96
8. Logging 97
8.1. Configuring Local Syslog 97
8.2. Configuring the Remote Syslog Server 97
8.3. Deleting Logs 100
9. SNMP 101
9.1. SNMP Traps 101
9.2. SNMP Access Configuration 103
9.2.1. Add an SNMP User ID 103
9.2.2. Create an SNMP Community 104
9.2.3. Map the Community to a Security Group 105
9.3. SNMP menu 105
9.4. SNMP Discovery 109
9.5. SNMP Community 109
9.6. SNMP Target Addresses 110
9.7. SNMP Users 112
9.8. SNMP Security to Group Maps 114
9.9. SNMP Access 114
10. Authentication 117
10.1. RADIUS 117
10.1.1. RADIUS overview 117
10.1.2. RADIUS Usage 117

10.1.3. RADIUS on ROX™ 118
10.1.4. RADIUS, ROX™, and Services 118
10.1.5. RADIUS Authentication Configuration 118
11. NETCONF 121
12. Chassis Management 125
12.1. Power Controller 126
12.2. Slot Hardware 127
12.3. Slot Identification 128
12.4. CPU 129
12.5. Slot Status 130
12.6. Slot Sensors 131
12.7. Module Configuration 132
13. PPP Users 135
13.1. Overview 135
13.2. PPP Configuration 135
13.3. PPP Interfaces and Link Failover 138
14. DHCP Relay 140
15. DHCP Server 142
15.1. DHCP Fundamentals 142
15.1.1. DHCP Network Organizations 142
15.1.2. Option 82 Support with Disable NAK 142
15.2. Configuring DHCP Server 143
15.2.1. Enabling the DHCP Service 143
15.2.2. DHCP Interfaces 143
15.2.3. DHCP Subnets and Pools 144
15.2.4. DHCP Shared Networks 145
15.2.5. DHCP Hosts 145
15.2.6. DHCP Host-groups 146
15.2.7. Viewing Active DHCP Leases 146
15.2.8. DHCP Options 147
15.2.9. Custom DHCP Options 152
15.2.10. Hardware Configuration 152
II. Network Interfaces and Ethernet Bridging 154
16. Ethernet Ports 155
16.1. Controller Protection Through Link-Fault-Indication (LFI) 155
16.2. Ethernet Port Configuration 156
16.2.1. Port Parameters 157
16.2.2. Port Rate Limiting 159
16.2.3. Port Mirroring 160
16.2.4. Diagnostics 162
16.2.5. Link Detection Options 167
16.3. Port Status 168
16.4. Resetting Ports 170
16.4.1. Resetting All Switched Ports 171
16.5. Troubleshooting 171
17. Ethernet Statistics 173
17.1. Viewing Ethernet Statistics 173
17.2. Viewing Ethernet Port Statistics 173
17.3. Viewing Non-switched Ethernet Statistics 178
17.4. Clearing Switched Ethernet Port Statistics 181
18. IP Statistics 183
19. Virtual Switch Bridging 186
19.1. Overview 186
19.1.1. Helpful Hints 186

19.2. Sample Use Case 187
19.3. Virtual Switch Configuration and Status 188
20. Link Aggregation 194
20.1. Link Aggregation Operation 194
20.1.1. Link Aggregation Rules 194
20.1.2. Link Aggregation Limitations 195
20.2. Link Aggregation Configuration 196
20.2.1. Configuring Port Trunks 196
21. Modem 203
21.1. PPP and the Cellular Modem 203
21.1.1. PPP and Cellular Modem Fundamentals 203
21.1.2. PPP Cellular Modem Information and Configuration 203
22. Serial Protocols 218
22.1. Introduction 218
22.1.1. Serial IP Port Features 218
22.1.2. Serial Protocols Applications 219
22.1.3. Serial Protocols Concepts And Issues 220
22.1.4. TcpModBus Server Application 221
22.1.5. TcpModbus Concepts And Issues 221
22.1.6. DNP (Distributed Network Protocol) 224
22.2. Serial Protocol Configuration 225
22.2.1. Assigning Protocols 225
22.2.2. Setting Rawsockets 228
22.2.3. Setting TcpModbus 229
22.2.4. Setting DNP 231
22.3. Serial Protocol Statistics 232
22.3.1. Transport Connections 234
22.4. Restarting the Serial Server 236
22.5. Resetting Ports 236
23. WAN 237
23.1. T1/E1 Fundamentals 237
23.1.1. Frame Relay 237
23.1.2. RX1500 and Frame Relay Encapsulation 237
23.2. WAN Configuration 238
23.2.1. T1 Parameters 239
23.2.2. E1 Parameters 240
23.2.3. Configuring Protocols 240
23.2.4. Loopback Test 248
23.3. Statistics 249
23.3.1. Physical Layer-related Statistics 250
23.3.2. Protocol-related Statistics 255
23.3.3. Clearing Statistics 261
23.4. DDS 261
23.4.1. DDS Configuration 262
23.4.2. Viewing and Clearing DDS Statistics 266
24. Port Security 268
24.1. Port Security Operation 268
24.1.1. Static MAC address-based authorization 268
24.1.2. IEEE 802.1X Authentication 268
24.2. Port Security Configuration 270
24.2.1. Port Security Parameters 271
24.2.2. 802.1X Parameters 272
25. Multicast Filtering 274
25.1. IGMP 274

25.1.1. Router and Host IGMP Operation 274
25.1.2. Switch IGMP Operation 275
25.1.3. Combined Router and Switch IGMP Operation 277
25.2. GMRP (GARP Multicast Registration Protocol) 277
25.2.1. GMRP Example 278
25.3. Multicast Filtering Configuration and Status 280
25.3.1. Configuring IGMP Parameters 280
25.3.2. Configuring Static Multicast Groups 282
25.3.3. Configuring GMRP 285
25.4. Troubleshooting 287
26. Classes Of Service 289
26.1. CoS Operation 289
26.1.1. Inspection Phase 289
26.1.2. Forwarding Phase 290
26.2. CoS Configuration 290
26.2.1. Global CoS Parameters 290
26.2.2. Priority to CoS Mapping 291
26.2.3. DSCP to CoS Mapping 292
27. MAC Address Tables 294
28. Spanning Tree 298
28.1. RSTP Operation 298
28.1.1. RSTP States and Roles 299
28.1.2. Edge Ports 301
28.1.3. Point-to-Point and Multipoint Links 301
28.1.4. Path and Port Costs 301
28.1.5. Bridge Diameter 302
28.2. MSTP Operation 302
28.2.1. MST Regions and Interoperability 303
28.2.2. MSTP Bridge and Port Roles 304
28.2.3. Benefits of MSTP 305
28.2.4. Implementing MSTP on a Bridged Network 305
28.3. RSTP Applications 306
28.3.1. RSTP in Structured Wiring Configurations 306
28.3.2. RSTP in Ring Backbone Configurations 308
28.3.3. RSTP Port Redundancy 309
28.4. Spanning Tree Configuration 309
28.4.1. Spanning Tree Parameters 310
28.4.2. Port RSTP Parameters 314
28.4.3. Bridge MSTI Parameters 316
28.4.4. Port MSTI Parameters 318
28.5. Spanning Tree Statistics 320
28.5.1. Bridge RSTP Statistics 320
28.5.2. Port RSTP Statistics 322
28.5.3. MSTI Status 325
28.5.4. Port MSTP Statistics 327
28.6. Clearing Spanning Tree Statistics 328
28.7. Troubleshooting 329
29. Virtual LANs 332
29.1. VLAN Operation 332
29.1.1. VLANs and Tags 332
29.1.2. Tagged vs. Untagged Frames 332
29.1.3. Native VLAN 332
29.1.4. Edge and Trunk Port Types 332
29.1.5. VLAN Ingress and Egress Rules 333

29.1.6. Forbidden Ports List 333
29.1.7. VLAN-aware Mode of Operation 333
29.1.8. GVRP (GARP VLAN Registration Protocol) 334
29.1.9. PVLAN Edge 335
29.2. VLAN Applications 336
29.2.1. Traffic Domain Isolation 336
29.2.2. Administrative Convenience 336
29.2.3. Reduced Hardware 336
29.3. VLAN Configuration 337
29.3.1. Static VLANs 338
29.3.2. Port VLAN Parameters 339
29.3.3. VLAN Summary 340
29.3.4. Forbidden Ports 343
29.4. Troubleshooting 343
30. Network Discovery 345
30.1. LLDP Operation 345
30.2. LLDP Parameters 346
III. Routing and Security 353
31. ROX™ Routing Overview 354
31.1. IP Routing in ROX™ 354
31.2. Physical Ethernet Port Types in ROX™ 354
31.3. Routing 354
31.3.1. Using VLAN Interfaces for Routing and Layer 3 Switching 354
31.3.2. Routing IP Packets 355
32. Layer 3 Switching 356
32.1. Layer 3 Switching Fundamentals 356
32.1.1. What is a Layer 3 Switch? 356
32.1.2. Layer 3 Switch Forwarding table 356
32.1.3. Static Layer 3 Switching Rules 357
32.1.4. Dynamic Learning of Layer 3 Switching Rules 357
32.1.5. Layer 3 Switch ARP table 357
32.1.6. Layer 3 Multicast Switching 358
32.1.7. Size of the Layer 3 Switch Forwarding Table 358
32.1.8. Interaction with the Firewall 358
32.1.9. Sample Use Case 359
32.2. Configuring Layer 3 Switching 362
32.2.1. Configuring Layer 3 Switching Settings 363
32.2.2. Creating Static ARP Table Entries 364
32.2.3. Viewing Static and Dynamic ARP Table Entries 365
32.2.4. Viewing Routing Rules 365
32.2.5. Flushing Dynamic Hardware Routing Rules 368
33. Tunnelling 369
33.1. IPsec 369
33.1.1. VPN Fundamentals 369
33.1.2. IPsec Configuration 372
33.2. L2TP Tunnelling Configuration 382
33.3. Layer 2 Tunnelling 384
33.3.1. IEC61850 GOOSE Fundamentals 384
33.3.2. Generic Layer 2 Tunnel Fundamentals 385
33.3.3. Layer 2 Tunnelling Configuration 386
33.4. Generic Routing Encapsulation (GRE) 394
33.4.1. Generic Routing Encapsulation Configuration 394
34. Dynamic Routing 397
34.1. Introduction 397

34.1.1. RIP, OSPF, and BGP 397
34.1.2. RIP Fundamentals 397
34.1.3. OSPF Fundamentals 397
34.1.4. Key OSPF And RIP Parameters 398
34.1.5. OSPF And VRRP Example Network 400
34.1.6. BGP Fundamentals 402
34.2. Dynamic Routing Configuration 402
34.3. RIP 402
34.3.1. RIP Configuration 403
34.4. OSPF 408
34.4.1. OSPF Configuration 409
34.5. BGP 413
34.5.1. BGP configuration 413
35. Static Routing 420
36. Routing Status 422
36.1. IPv4 422
36.2. IPv6 423
36.3. Memory Statistics 423
36.4. RIP 425
36.5. OSPF 426
36.6. BGP 430
37. Multicast Routing 433
38. Firewall 437
38.1. Firewall Fundamentals 437
38.1.1. Stateless vs Stateful Firewalls 437
38.1.2. Linux® netfilter, iptables, and the Firewall 437
38.1.3. Network Address Translation 437
38.1.4. Port Forwarding 438
38.2. Firewall Quick Setup 438
38.3. Firewall Terminology And Concepts 439
38.3.1. Zones 439
38.3.2. Interfaces 439
38.3.3. Hosts 440
38.3.4. Policy 440
38.3.5. Masquerading and SNAT 441
38.3.6. Rules 442
38.4. Configuring The Firewall And VPN 443
38.4.1. Policy-based Virtual Private Networking 443
38.4.2. Virtual Private Networking to a DMZ 444
38.5. Firewall Configuration 444
38.5.1. Adding a Firewall 445
38.5.2. Working with Firewall Configurations 446
38.5.3. Zone Configuration 447
38.5.4. Interface Configuration 448
38.5.5. Host Configuration 449
38.5.6. Policies 450
38.5.7. Network Address Translation 451
38.5.8. IP Masquerading 452
38.5.9. Rules 453
39. Traffic Control 457
39.1. Traffic Control Modes 457
39.1.1. Traffic Control Basic (basic-configuration) Configuration Mode 457
39.1.2. Traffic Control Advanced (advanced-configuration) Configuration Mode 457

39.2. Traffic Control Configuration 459
39.2.1. Traffic Control Modes 459
40. VRRP 476
40.1. VRRP Fundamentals 476
40.1.1. The Problem With Static Routing 476
40.1.2. The VRRP Solution 476
40.1.3. VRRP Terminology 476
40.2. VRRP Configuration 478
40.2.1. VRRP Status 481
41. Link Failover 483
41.1. Path Failure Discovery 483
41.2. Using Routing Protocols and the Default Route 483
41.3. Configuring Link Failover 483
41.3.1. Configuring the Link Failover Settings 484
41.3.2. Setting a Link Failover Backup Interface 485
41.3.3. Setting a Link Failover Ping Target 486
41.3.4. Link Backup On Demand 487
41.3.5. Viewing Link Failover Status 487
41.3.6. Viewing the Link Failover Log 488
41.3.7. Testing Link Failover 488
IV. Appendices 490
A. Upgrading Software 491
A.1. Preparing The Software Upgrade 491
A.2. Launching The Upgrade 492
A.3. Monitoring The Software Upgrade 493
B. RADIUS Server Configuration 497
B.1. PPP / CHAP and Windows IAS 497
C. Setting Up An Upgrade Server 498
C.1. Upgrade Server Requirements 498
C.2. Initial Upgrade Server Setup 498
C.3. Upgrading The Repository 498
C.4. Setting Up The Routers 499
C.5. Using Microsoft Internet Information Services (IIS) Manager 6.0 or Higher as a ROX Upgrade Repository 499
D. Adding and Replacing Line Modules 500
D.1. Shutting Down the Unit to Remove/Insert Modules 500
D.2. Adding a Module to an Empty Slot 500
D.3. Swapping a Module with an Identical Backup Module 500
D.4. Swapping a Module with a Different Type of Module 500
D.5. Swapping a Module with a Different Type of Module 501
E. GNU General Public License 502
E.1. Preamble 502
E.2. TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION 503
E.2.1. Section 0 503
E.2.2. Section 1 503
E.2.3. Section 2 503
E.2.4. Section 3 504
E.2.5. Section 4 504
E.2.6. Section 5 504
E.2.7. Section 6 504
E.2.8. Section 7 505
E.2.9. Section 8 505
E.2.10. Section 9 505

E.2.11. Section 10 505
E.2.12. NO WARRANTY Section 11 506
E.2.13. Section 12 506
E.3. How to Apply These Terms to Your New Programs 506

List of Figures

1.1. The ROX™ Login page 28
1.2. The ROX™ Web Interface 28
1.3. Top-level Menu 30
1.4. Example of Edit Private Mode 32
1.5. Adding Key Information 33
1.6. Key Information in a Table 34
1.7. Example of Key Settings 1 34
1.8. Example of Key Settings 2 35
1.9. First Table of Information 36
1.10. Second Table of Information 36
2.1. Administration menu 37
2.2. Clear All Alarms Menu Action form 37
2.3. Acknowledge All Alarms Menu Action form 37
2.4. Shutdown the Device Menu Action form 38
2.5. Reboot the Device Menu Action form 38
2.6. Set New Time and Date form 38
2.7. Set Clock on Target Device form 38
2.8. Restore-factory-defaults Trigger Action form 39
2.9. Administration form 39
2.10. Hostname form 39
2.11. Timezone form 40
2.12. Setting the Timezone Form - in Edit Private Mode 40
2.13. Current System Time form 40
2.14. CLI Sessions form 41
2.15. Idle-timeout field 42
2.16. Session Limits form 42
2.17. SFTP Sessions form 43
2.18. WWW Interface Sessions 44
2.19. Idle-timeout field 45
2.20. Users menu 45
2.21. Users table 45
2.22. Users form 46
2.23. Users Screen in Edit Private View 46
2.24. Software-Upgrade menu 47
2.25. Upgrade Settings 47
2.26. Upgrade Monitoring 48
2.27. Launch Upgrade 49
2.28. Decline Upgrade 49
2.29. Rollback and Reboot 49
2.30. ROX-Imaging menu 50
2.31. ROXflash Monitoring form 51
2.32. ROXFlash menu 51
2.33. ROXFlash forms 52
2.34. Scheduler menu 52
2.35. Scheduled-jobs table 53
2.36. Scheduled Jobs Form 53
2.37. CLI in the ROX™ Web Interface 56
2.38. Install Files forms 57
2.39. Backup Files forms 59
2.40. Administration menu 59
2.41. Install Files forms 60

2.42. Backup Files forms 60
2.43. Delete-logs menu 61
2.44. Delete Log Files form 61
2.45. Save-full-configuration menu 61
2.46. Save Full Configuration forms 62
2.47. Load-full-configuration menu 62
2.48. Load Full Configuration forms 62
3.1. Set new Time and Date form 64
3.2. Timezone form 65
3.3. Local Time Settings form 65
3.4. Network Time Protocol (NTP) Servers 65
3.5. Network Time Protocol (NTP) Servers form 66
3.6. Server Keys form 67
3.7. Server Restrictions Key settings form 68
3.8. Server Restrictions form 68
3.9. NTP Broadcast/Multicast Servers form 69
3.10. NTP Multicast Clients form 70
3.11. Network Time Protocol (NTP) form 70
3.12. NTP Service Status form 71
4.1. IP menu 72
4.2. Configuring an IP Address 73
4.3. Basic Network Setup Using the Default IPv4 Addresses 74
4.4. Simple IPv6 Network Setup 75
4.5. Routable Interfaces table 76
4.6. Routable Interfaces form 76
4.7. Addresses table 76
4.8. Addresses form 76
5.1. Neighbor Discovery form 79
5.2. Neighbor Discovery IPv6 Prefix 80
5.3. Neighbor Discovery IPv6 Prefix forms 80
5.4. Explicitly Adding a VLAN Interface to a Switched Port 82
5.5. All VLANs table 84
5.6. All VLANs Properties form 84
5.7. Non-switched Interface menu 85
5.8. Routable Ethernet Ports table 85
5.9. Routable Ethernet Ports form 85
5.10. Configuring Dynamic Address Source and ProxyARP 87
6.1. Alarms menu 90
6.2. Active Alarms table 90
6.3. Active Alarms Key Settings form 90
6.4. Active Alarms form 91
6.5. Clear Alarm Menu Action form 92
6.6. Acknowledge Alarm Menu Action form 92
6.7. Clear All Alarms Menu Action form 92
6.8. Acknowledge All Alarms Menu Action form 92
6.9. Admin Alarm Configuration table 93
6.10. Admin Alarm Configuration form 93
6.11. Chassis Alarm Configuration table 94
6.12. Chassis Alarm Configuration form 94
6.13. Switch Alarm Configuration table 95
6.14. Switch Alarm Configuration form 95
7.1. DNS menu 96
7.2. Domain Name Searches form 96
7.3. Domain Name Servers 96

8.1. Logging menu 97
8.2. Remote Server table 97
8.3. Remote Server form 98
8.4. Remote Server Selector table 98
8.5. Selector menu 98
8.6. Remote Server Selector form 99
9.1. Adding an SNMP User ID 103
9.2. Creating an SNMP Community 104
9.3. Mapping the Community to a Security Group 105
9.4. SNMP menu 105
9.5. SNMP Sessions form 106
9.6. SNMP USM Statistics form 108
9.7. SNMP-Discover action 109
9.8. SNMP Engine ID Discover forms 109
9.9. SNMPv1/v2c Community Configuration table 109
9.10. SNMPv1/v2c Community Configuration form 110
9.11. SNMP Target Configuration table 110
9.12. SNMPv3 Target Configuration form 111
9.13. SNMP User Configuration table 112
9.14. User Configuration Key Settings form 113
9.15. SNMP User Configuration form 113
9.16. SNMP Security Model to Group Mapping table 114
9.17. Key Settings form 114
9.18. SNMP Security Model to Group Mapping form 114
9.19. SNMP Group Access Configuration table 115
9.20. Key Settings form 115
9.21. SNMP Group Access Configuration form 115
10.1. Authentication menu 117
10.2. Primary RADIUS Server form 119
10.3. Secondary RADIUS Server form 119
11.1. NETCONF menu 121
11.2. NETCONF Sessions form 121
11.3. Idle-timeout field 122
11.4. NETCONF State/Statistics form 123
12.1. Chassis menu 125
12.2. Chassis Status form 125
12.3. Power Controller form 126
12.4. Power Status table 126
12.5. Power Status form 126
12.6. Slot Hardware table 127
12.7. Slot Hardware form 127
12.8. Slot Identification table 128
12.9. Slot Identification form 128
12.10. Slot CPU/RAM Utilization table 129
12.11. Slot CPU/RAM Utilization form 129
12.12. Slot Status table 130
12.13. Slot Status form 130
12.14. Slot Sensors table 131
12.15. Slot Sensors form 131
12.16. Modules table 132
12.17. Modules form 132
12.18. Fixed Modules form 133
12.19. Fixed Modules table 133
12.20. Module Database table 134

12.21. Module Database form 134
12.22. Configurable Modules table 134
12.23. Configurable Modules form 134
13.1. PPP menu 135
13.2. Dial-in PPP Users table 135
13.3. Dial-in Users form 136
13.4. Dial-out PPP Users table 136
13.5. PPP Configuration form 136
13.6. PPP Primary Radius Server form 138
13.7. PPP Secondary Radius Server form 138
14.1. DHCP Relay Agent Menu 140
14.2. DHCP Relay Agent Form 140
14.3. DHCP Relay Agent Client Ports table 141
15.1. DHCP Server menu 143
15.2. DHCP Server form 143
15.3. Listen Interfaces table 144
15.4. Subnet Configuration table 144
15.5. Subnet Configuration form 144
15.6. IP Pool Configuration table 145
15.7. Shared Network Configuration table 145
15.8. Host Configuration table 145
15.9. Host Group Configuration table 146
15.10. /services/dhcpserver/show-active-leases form 147
15.11. Lease Configuration form 148
15.12. Client Configuration form for Subnets and Shared Networks 148
15.13. Client Configuration form for Hosts 149
15.14. Client Configuration form for Host-groups 149
15.15. Client Configuration form for DHCP Clients 150
15.16. NIS Configuration form 151
15.17. Netbios Configuration form 151
15.18. Setting a DHCP Custom Option 152
15.19. Hardware Configuration form 152
16.1. Controller Protection Through LFI 155
16.2. Ethernet Ports menu 156
16.3. Switched Ethernet Ports table 157
16.4. Switched Ethernet Ports submenu 157
16.5. Switched Ethernet Ports form 158
16.6. Rate Limiting form 159
16.7. Port-Mirroring menu 161
16.8. Port Mirror form 161
16.9. Ingress Source Ports table 161
16.10. Egress Source Ports table 161
16.11. Diagnostics menu 162
16.12. Cable Diagnostics Results form 162
16.13. Start Cable Diagnostics Test form 165
16.14. Start Cable Test form 165
16.15. Clear Port Cable Diagnostic Test Results form 165
16.16. Clear All Diagnostics (Switch) menu 166
16.17. Clear All Cable Diagnostic Test Results form 166
16.18. Clear All Alarms menu 166
16.19. Clear All Active Alarms Trigger Action 166
16.20. Switch (Link Detection) menu 167
16.21. Link Detection form 167
16.22. Interfaces menu 168

16.23. Interface Status table 169
16.24. Interface Status form 169
16.25. Port Security Status form 170
16.26. Reset Ethernet Port form 171
16.27. Reset All Switched Ports menu 171
16.28. Reset All Switched Ports form 171
17.1. Ethernet Port Statistics Menu 173
17.2. Port Statistics Form 173
17.3. RMON Port Statistics Form 175
17.4. Statistics Menu 178
17.5. Routable-Only Ethernet Port Status Form 179
17.6. Receive Statistics Form 180
17.7. Transmit Statistics Form 181
17.8. Interfaces Switch (Clearing Port Statistics) Menu 181
17.9. Clear Switched Port Statistics Form 182
17.10. Clear All Statistics Menu 182
17.11. Clear All Switched Port Statistics Form 182
18.1. Interfaces IP Menu 183
18.2. Routable Interface Statistics Table 183
18.3. Routable Interface Statistics Form 183
18.4. Receive Statistics Form 184
18.5. Transmit Statistics Form 184
19.1. Virtual switch with multiple interfaces 187
19.2. Adding a Virtual Switch 188
19.3. Interface Virtualswitch menu 188
19.4. Virtualswitch table 188
19.5. Virtualswitch form 189
19.6. Interface table 189
19.7. VLAN table 189
19.8. VLAN form 190
19.9. Interfaces Virtualswitch menu 190
19.10. Virtualswitch table 190
19.11. Virtualswitch form 190
19.12. Receive form 191
19.13. Transmit form 191
19.14. VLAN table 192
19.15. VLAN Receive form 192
19.16. VLAN Transmit form 193
20.1. Link Aggregation Examples 194
20.2. Link Aggregation menu 196
20.3. Adding Trunks 196
20.4. Entering a Trunk ID 197
20.5. Entering Parameters for Forms 198
20.6. Trunk-Ports Submenu - Adding a Trunk-Port 199
20.7. Selecting a Trunk Slot 199
20.8. Trunk Ports table 200
20.9. Trunk Ports Table in Edit Private Mode 200
20.10. Key Settings 200
20.11. Ethernet Trunk Interfaces form 200
20.12. Multicast Filtering form 200
20.13. CoS form 201
20.14. VLAN form 201
20.15. Trunk Ports table 202
21.1. Interfaces Cellmodem menu 203

21.2. HSPA Cellular Modem Information form 204
21.3. Edge Cellular Modem Information form 205
21.4. Global Cellular GSM menu 206
21.5. GSM Cellular Network Configuration form 207
21.6. PPP Configuration form 207
21.7. CDMA EVDO Cellular Modem Information form 208
21.8. CDMA Over The Air Activation form 210
21.9. CDMA Over The Air Activation Trigger Action form 210
21.10. CDMA Manual Activation form 211
21.11. CDMA Manual Activation Trigger Action form 211
21.12. CDMA Reset Modem Trigger Action form 211
21.13. Global Cellular CDMA menu 212
21.14. Cellular Network Configuration table 212
21.15. Cellular Network Configuration form 212
21.16. PPP Configuration form 213
21.17. Interface Cellmodem menu 213
21.18. Routable Cellular Modem Interfaces table 214
21.19. Routable Cellular Modem Interfaces form 214
21.20. Interface Cellmodem HSPA menu 215
21.21. GSM Profile form 215
21.22. Interfaces Cellmodem menu 215
21.23. Cellular Modem Interfaces table 215
21.24. Interfaces Cellmodem HSPA menu 216
21.25. Cellular Modem Interfaces form 216
21.26. HSPA PPP Interfaces Statistics form 216
22.1. 6S01 Serial Module RJ45 Connector LEDs 218
22.2. Sources of Delay and Error in an End to End Exchange 222
22.3. Serial Protocols menu 225
22.4. Serial Interfaces table 225
22.5. Adding a Protocol in the Edit Private screen 226
22.6. Selecting a Protocol Type in the Edit Private screen 226
22.7. Serial Ports Configuration form 227
22.8. Serial Protocols table 228
22.9. Rawsocket Configuration form 228
22.10. TCP Modbus Configuration form 229
22.11. DNP Protocols Configuration form 231
22.12. DNP Device Address Table Configuration table 231
22.13. DNP Device Address Table Configuration form 231
22.14. Serial Protocol Statistics menu 232
22.15. Serial Port Statistics table 232
22.16. Serial Port Statistics form 233
22.17. Transport Connections Statistics table 234
22.18. TCP/UDP Connection Statistics form 235
22.19. Restart Serserver menu 236
22.20. Restart Serserver Trigger Action 236
22.21. Reset Ports menu 236
22.22. Reset Ports Trigger Action 236
23.1. WAN menu 238
23.2. Interface WAN Slot Port Settings table 238
23.3. Enable WAN Interface form 238
23.4. T1 Parameters form 239
23.5. E1 Parameters form 240
23.6. T1 Channels and Associated Time Slots table 241
23.7. T1 Time Slots form 241

23.8. Adding a Connection 242
23.9. Frame Relay Parameter form 242
23.10. Connection Frame Relay DLCI table 243
23.11. Adding an MLPPP Connection 244
23.12. Adding IP and Remote Addresses 245
23.13. HDLC-ETH menu 246
23.14. Ethernet Over HDLC Settings form 246
23.15. Adding a VLAN 247
23.16. T1/E1 Interfaces under the IP submenu 248
23.17. Loopback Test Forms 248
23.18. Loopbacktest Results 249
23.19. WAN Statistics menu 249
23.20. T1E1 Statistics table 249
23.21. Receiving Errors Statistics form 250
23.22. T1E1 Receiving Statistics form 251
23.23. T1E1 Receiving Statistics Form 2 251
23.24. T1E1 Transmitting Errors Statistics form 252
23.25. T1E1 Transmitting Statistics form 252
23.26. T1E1 Transmitting Statistics Form 2 253
23.27. T1E1 Alarm Indication form 254
23.28. T1E1 Statistics form 255
23.29. PPP Receiving Protocol Statistics form 255
23.30. PPP Transmitting Protocol Statistics form 256
23.31. T1E1 Statistics form 256
23.32. Frame Relay Errors Packets Statistics form 258
23.33. Frame Relay Controlling Packets Statistics form 259
23.34. Frame Relay Receiving Statistics form 260
23.35. Clear Interface Statistics Form And Trigger Action 261
23.36. Clearstatistics Menu Action 261
23.37. Enable Wan Interface form 262
23.38. DDS Parameters form 263
23.39. PPP form 263
23.40. Frame Relay Parameters form 264
23.41. Loopback Test form 265
23.42. DDS Statistics menu 266
23.43. Clear Interface Statistics form 267
24.1. 802.1X General Topology 269
24.2. 802.1X Packet Exchange 269
24.3. Port Security RADIUS Primary form 270
24.4. Port Security RADIUS Secondary form 270
24.5. Port Security menu 271
24.6. Port Security form 271
24.7. 802.1x Parameters form 272
25.1. IGMP Operation Example 1 275
25.2. IGMP Operation Example 2 277
25.3. Example using GMRP 279
25.4. Multicast Filtering menu 280
25.5. IGMP Snooping Parameters form 281
25.6. Router Ports table 281
25.7. Egress Ports table 282
25.8. Static Multicast Summary table 282
25.9. Static Multicast Summary form 282
25.10. Static Ports table 283
25.11. Static Ports form 283

25.12. Multicast Group Summary table 283
25.13. IP Multicast Groups table 284
25.14. IP Multicast Groups form 284
25.15. Router-Ports table 284
25.16. Router-Ports form 284
25.17. Joined-Ports table 285
25.18. Joined-Ports form 285
25.19. GMRP form 285
25.20. GMRP Dynamic Ports table 286
25.21. GMRP Dynamic Ports form 286
25.22. Multicast Filtering form 286
26.1. Determining The CoS Of A Received Frame 290
26.2. Class-of-service menu 290
26.3. CoS form 290
26.4. Priority to CoS Mapping table 291
26.5. Priority to CoS Mapping form 291
26.6. TOS DSCP to CoS Mapping table 292
26.7. TOS DSCP to CoS Mapping form 292
26.8. CoS form 292
27.1. MAC Tables menu 294
27.2. MAC Address table 294
27.3. Mac Address form 294
27.4. MAC Tables form 295
27.5. Key Settings 296
27.6. Static MAC Address Parameters form 296
27.7. Static MAC Address Parameters table 296
27.8. Purge MAC Address menu 297
27.9. Purge MAC Address Table form 297
28.1. Bridge and Port States 299
28.2. Bridge and Port Roles 300
28.3. Example of a Structured Wiring Configuration 307
28.4. Example of a Ring Backbone Configuration 308
28.5. Port Redundancy 309
28.6. Spanning Tree menu 310
28.7. Spanning Tree Parameter form 310
28.8. RSTP Common Instance form 312
28.9. eRSTP form 312
28.10. Interface/switch/{line module}/spanning-tree submenu 314
28.11. Port RSTP Parameter form 314
28.12. Key Settings form 316
28.13. MSTP Instance form 316
28.14. MSTP Instance table 317
28.15. MSTP ID table 317
28.16. MSTI Configuration table 318
28.17. MSTI Configuration form 318
28.18. RSTP Status form 320
28.19. RSTP Port Statistics table 322
28.20. RSTP Port Statistics form 323
28.21. MSTI Status table 325
28.22. MSTI Status form 325
28.23. MSTP Port Statistics table 327
28.24. MSTP Port Statistics form 327
28.25. Clear-stp-stats Menu Action 328
28.26. Clear Spanning-Tree Statistics form 329

29.1. Using GVRP 335
29.2. Multiple Overlapping VLANs 336
29.3. Inter-VLAN Communications 337
29.4. Virtual LANs menu 337
29.5. Internal VLAN Range form 338
29.6. Static VLAN table 338
29.7. Static VLAN form 338
29.8. Switched Ethernet Ports submenu 339
29.9. VLAN Parameters form 339
29.10. VLAN Summary table 340
29.11. VLAN Summary form 341
29.12. Tagged Ports table 341
29.13. Tagged Ports form 341
29.14. Untagged Ports table 342
29.15. Untagged Ports form 342
29.16. All VLANs table 342
29.17. All VLANs Properties form 342
29.18. VLANs table 342
29.19. VLANs form 343
29.20. Forbidden Ports 343
30.1. Net-discovery menu 345
30.2. Net-discovery LLDP menu 346
30.3. LLDP form 346
30.4. LLDP Global Statistics form 347
30.5. LLDP Local System form 348
30.6. LLDP Port Statistics table 349
30.7. LLDP Port Statistics form 349
30.8. LLDP Neighbors table 350
30.9. LLDP Neighbors form 351
30.10. LLDP submenu 351
30.11. LLDP form 352
31.1. Three interfaces on an isolated VLAN 354
31.2. VLAN connected to ROX device through switch.0100 355
32.1. Layer 3 Switch 356
32.2. Layer 3 Switch Use Case 359
32.3. Hardware Acceleration Enabled 360
32.4. Hardware Acceleration Enabled 360
32.5. Layer 3 Switching menu 362
32.6. Layer 3 Switching form 362
32.7. Layer 3 Switching form 363
32.8. ARP Table Configuration form 365
32.9. ARP Table Summary form 365
32.10. Routing Rules Summary table 366
32.11. Routing Rules Summary form 366
32.12. Flush Dynamic Hardware Routing Rules form 368
33.1. Tunnelling menu 369
33.2. IPsec menu 372
33.3. IPsec form 372
33.4. Syslog form 372
33.5. Show Public RSA Key form 373
33.6. Install-Certificate forms 374
33.7. Install-Ca-Certificate forms 375
33.8. Install-Crl-File forms 376
33.9. Show IPsec Running Status form 376

33.10. Connection table 376
33.11. Connection form 377
33.12. ESP table 378
33.13. ESP Key Settings 378
33.14. IKE table 378
33.15. Public IP Address form 379
33.16. System Public Key form 379
33.17. Nexthop To Other System form 380
33.18. System Identifier form 380
33.19. Private Subnet Behind System form 380
33.20. Network table 381
33.21. Preshared Key table 381
33.22. Preshared Key form 381
33.23. L2TP menu 382
33.24. L2TP form 382
33.25. DNS Server form 382
33.26. PPP Options form 383
33.27. WINS Server form 383
33.28. L2tunneld menu 386
33.29. L2 Tunnel Daemon form 386
33.30. Goose Tunnel table 387
33.31. Goose Tunnel form 387
33.32. Remote Daemon of Goose Tunnel table 387
33.33. Generic L2 Tunnel table 387
33.34. Generic L2 Tunnel Protocol form 388
33.35. Generic L2 Tunnel Egress Interface table 388
33.36. L2 Ethernet Type table 388
33.37. Goose Tunnel Statistics table 388
33.38. Goose Tunnel Statistics form 389
33.39. Connections Statistics table 390
33.40. Connections Statistics form 390
33.41. Generic L2 Tunnel Statistics table 391
33.42. Generic L2 Tunnel Statistics form 391
33.43. Connections Statistics table 392
33.44. Connections Statistics form 392
33.45. Round Trip Time Statistics table 393
33.46. Round Trip Time Statistics form 393
33.47. GRE Example 394
33.48. Generic Routing Encapsulation (GRE) menu 394
33.49. Generic Routing Encapsulation Interfaces table 395
33.50. Generic Routing Encapsulation Interfaces form 395
34.1. OSPF and VRRP Example 401
34.2. Dynamic Routing Menu 402
34.3. RIP Menu 402
34.4. RIP Configuration Form 403
34.5. Routing Timers Form 404
34.6. RIP Interface Parameters Table 406
34.7. RIP Interface Parameters Form 407
34.8. Authentication Form 407
34.9. OSPF Menu 408
34.10. OSPF Configuration Form 409
34.11. OSPF Area Distance Form 410
34.12. Interface Parameters Table 411
34.13. Interface Parameters Form 411

34.14. Dead Interval Form 412
34.15. BGP Menu 413
34.16. BGP Configuration Form 413
34.17. Distance Form 414
35.1. Static Menu 420
35.2. Static Route table 420
35.3. Static Route form 420
35.4. Static Route Using Gateway table 420
35.5. Static Route Using Gateway form 420
35.6. Blackhole Static Route form 421
35.7. Static Route Using Interface table 421
35.8. Static Route Using Interface form 421
36.1. Routing Status Menu 422
36.2. IPv4 Kernel Active Routing Table 422
36.3. IPv6Kernel Active Routing Table 423
36.4. Core Daemon Memory Statistics Form 424
36.5. RIP Daemon Memory Statistics Form 424
36.6. BGP Daemon Memory Statistics Form 424
36.7. OSPF Daemon Memory Statistics Form 425
36.8. RIP Menu 425
36.9. OSPF Menu 426
36.10. Network Table 426
36.11. Reach Table 426
36.12. Router Table 427
36.13. Area Table 427
36.14. Net Table 427
36.15. Summary Table 428
36.16. ASBR-Summary Table 429
36.17. AS-External Table 429
36.18. Neighbor Table 430
36.19. BGP Menu 430
36.20. Route Table 431
36.21. Next Hop Table 431
36.22. BGP Neighbor Table 432
37.1. Multicast Routing menu 433
37.2. Static Multicast Routing Configuration form 433
37.3. Static menu 433
37.4. Multicast Groups Configuration table 433
37.5. Multicast Groups Configuration form 434
37.6. Outgoing Interfaces table 434
37.7. Multicast Routing Status table 435
37.8. Multicast Routing Status form 435
38.1. Security Menu 444
38.2. Firewall Description table 444
38.3. Firewall Description form 444
38.4. Adding a Firewall 445
38.5. Naming a Firewall 445
38.6. Firewall Submenus 446
38.7. Firewall Configuration form 446
38.8. Zone table 447
38.9. Zone form 447
38.10. Main Interface Settings table 448
38.11. Interface Options form 448
38.12. Broadcast Address form 449

38.13. Main Host Settings table 449
38.14. Main Host Settings form 449
38.15. Host Options form 450
38.16. Main Policy Settings table 450
38.17. Main Policy Settings form 450
38.18. Destination Zone form 451
38.19. Source Zone form 451
38.20. Net Address Translation Main Settings table 451
38.21. Net Address Translation Main Settings form 452
38.22. FWMasq table 452
38.23. Net Address Translation Main Settings form 453
38.24. Main Rule Settings table 453
38.25. Main Rule Settings form 454
38.26. Source Zone form 455
38.27. Destination Zone form 455
39.1. Traffic-Control menu 459
39.2. Traffic Control Configuration form 459
39.3. Enabling Basic-configuration Mode 460
39.4. Basic Traffic Control Interfaces table 460
39.5. Interface to Apply Traffic Control form 461
39.6. Basic Traffic Control Priorities table 462
39.7. Priorities form 462
39.8. Enabling Advanced-configuration Mode 464
39.9. Advanced Traffic Control Classes table 465
39.10. TC Classes form 465
39.11. Options form 467
39.12. Advanced Traffic Control Interfaces table 468
39.13. TC Devices form 469
39.14. TCrules menu 470
39.15. Advanced Traffic Control Rules table 470
39.16. TCrules form 471
39.17. Set form 473
39.18. Modify form 474
39.19. Save form 474
39.20. Restore form 474
39.21. Continue form 475
40.1. VRRP Example 477
40.2. VRRP Group Example 478
40.3. VRRP Menu 478
40.4. Virtual Router Redundancy Protocol (VRRP) Form 479
40.5. VRRP Group Table 479
40.6. VRRP Instance Table 479
40.7. VRRP Instance Form 480
40.8. Monitor Interface Form 481
40.9. VRIP IP Address Table 481
40.10. VRRP Status Table 481
40.11. VRRP Status Form 482
41.1. Link Backup Example 483
41.2. Link Fail Over Information Table 484
41.3. Link Fail Over Settings form 484
41.4. Backup Settings form 486
41.5. Link Fail Over Status form 487
41.6. Link Fail Over Logs form 488
41.7. Link Fail Over Test Settings form 489

A.1. The Software Upgrade Menu Interface 491
A.2. Entry Fields in Upgrade Settings Form 492
A.3. Pending Commit 492
A.4. Commit Succeeded 492
A.5. Launch Upgrade 493
A.6. Upgrade Launched Dialogs 493
A.7. Software-Upgrade Menu 493
A.8. Upgrade Monitoring Form in Reboot-pending Stage 494
A.9. Upgrade Monitoring Form Showing Successful Upgrade 495

Preface

This guide describes the web-based user interface for the ROX™ version 2.2 Operating System running on the RuggedBackbone™ RX1500 family of products.

Supported Platforms

ROX™2.2 is designed to work on RuggedCom's RuggedBackbone™ and RuggedRouter® hardware platforms. This ensures a consistent user experience when migrating from one product model in the family to another.

ROX™ currently supports the following RuggedCom networking platforms:

- RuggedBackbone™ RX5000 family of rugged, modular, Layer 3 switching multi-service hardware platforms.
- RuggedBackbone™ RX1500 family of rugged, modular, hot-swappable Layer 3 switching and routing platforms.
- RuggedRouter® RX1000 family of rugged Cyber-Security Appliances.

Who Should Use This User Guide

This guide is recommended for use by network technical support personnel who are familiar with the operation of networks. Others who might find the book useful are network and system planners, system programmers, and line technicians.

How This Guide Is Organized

Part I, “Administration”

This part covers the graphical user interface and overall management of the hardware chassis and operating system, including access control, logging, networking configuration, and time synchronization.

Part II, “Network Interfaces and Ethernet Bridging ”

This part covers the configuration and monitoring of the Ethernet bridging functions of the system, including Ethernet port setup, the Spanning Tree Protocol, and Virtual LANs.

Part III, “Routing and Security”

This part covers the configuration and monitoring of layer 3 routing and security functions, including OSPF, RIP, BGP, Multicast, and the Firewall.

Each part of this guide is organized into chapters that are typically devoted to one particular feature of the system.

Each chapter discusses mechanisms, protocols, or techniques specific to a particular feature. Many chapters include a general overview of the feature or protocol to be configured, providing some background into the feature and how it is used on the device. All chapters present the forms and fields in the web interface through which you configure the feature.All chapters present the CLI commands you use to configure the feature.

While every effort is made to ensure the accuracy and completeness of this guide, some web interface illustrations may not be exactly as shown.

Applicable Operating System Software Revision

This guide is applicable to ROX™ version 2.2.

Part I. Administration

This part describes the administration of a ROX™-based networking device:

The ROX Web Interface	Chapter 1, *The ROX™ Web Interface*
System Administration	Chapter 2, *System Administration*
Time Synchronization	Chapter 3, *Time Synchronization*
Basic Networking Configuration	Chapter 4, *Basic Network Configuration*
Advanced Networking Configuration	Chapter 5, *IP Network Interfaces*
Alarms	Chapter 6, *Alarms*
Domain Name Search	Chapter 7, *Domain Name Search*
Logging	Chapter 8, *Logging*
SNMP Configuration	Chapter 9, *SNMP*
Authentication	Chapter 10, *Authentication*
Notifications	Chapter 11, *NETCONF*
Physical Chassis Configuration	Chapter 12, *Chassis Management*
PPP User Profiles	Chapter 13, *PPP Users*
DHCP Relay	Chapter 14, *DHCP Relay*
DHCP Server	Chapter 15, *DHCP Server*

1. The ROX™ Web Interface

ROX™ features two primary user interfaces: a web-based interface and a command line interface (CLI). This user guide documents the usage and structure of the web-based user interface. For details of the CLI, please refer to the ROX™ Command Line Interface User Guide (in progress).

1.1. Getting Started

1.1.1. Requirements

Accessing the ROX™ web interface for the first time, prior to any system configuration, requires:

- A computer with an installed web browser capable of running JavaScript. ROX™ supports the following web browsers:
  - Microsoft® Internet Exporer 8.0 and higher
  - Mozilla Firefox
  - GNU Iceweasel
  - Google Chrome
- The computer must have a working Ethernet interface, which must be compatible with at least one of the port types on the RuggedBackbone™ as ordered.
- The ability to configure an IP address and netmask on the computer's Ethernet interface.

1.1.2. Connecting To The Web Interface

By default, the RuggedBackbone™ RX1500 has a different IP address and subnet configured for each of two distinct IP interfaces, each of which is mapped to one or more physical ports:

Interface Name	Location	IP Address/Mask
fe-cm-1	Front panel interface	192.168.1.2/24
All other Ethernet ports	LM and SM cards	192.168.0.2/24

Table 1.1. Default IP Address Configuration

In order to connect to the RX1500 using a web browser, configure the IP address of the web browser's system to fall within the subnet of the corresponding RX1500 interface. For example, if the web browser system is connected to the Ethernet interface on the RX1500 front panel:

- The web browser system's Ethernet interface must be configured with an IP address in the range: 192.168.1.3 to 192.168.1.254.
- The RX1500 is accessible to the web browser at the IP address: 192.168.1.2, the address of the fe-cm-1 network interface.

1.1.3. The Web Browser Connection

The ROX™ web server uses SSL (Secure Socket Layer) to encrypt data traffic exchanged with its clients (connections made via "https://"). This guarantees the privacy of communications between browser and server.

It can happen that upon connecting to the ROX™ web server, some new web browsers may report that they cannot verify the authenticity of the server's certificate against any of their known certificate authorities. This is expected, and it is safe to instruct the browser to accept the certificate offered by the ROX™ system. Once the browser is instructed to accept the certificate, all communications with the web server will be secure.

Start a web browser session and open a connection to the switch by entering a URL that specifies its IP address (https://192.168.1.2, to continue with the example above). After the web browser makes contact with the switch, the login page appears:

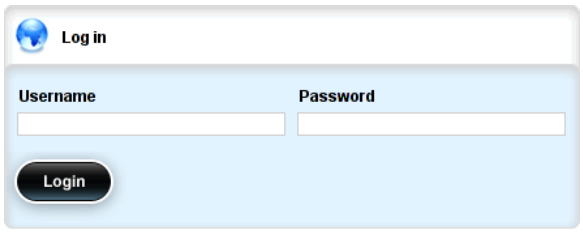

Figure 1.1. The ROX™ Login page

Enter the default user name, "admin" and the configured password for the admin user. Click on the **Login** button. The switch is shipped with a default administrator password, "admin". If authentication is successful, the main menu is presented.

1.2. The Structure Of The Web Interface

The system configuration interface (the *Configure Running* tab) is organized as a hierarchical set of linked menu entries, which may be traversed using the four-panel navigation window, as illustrated below.

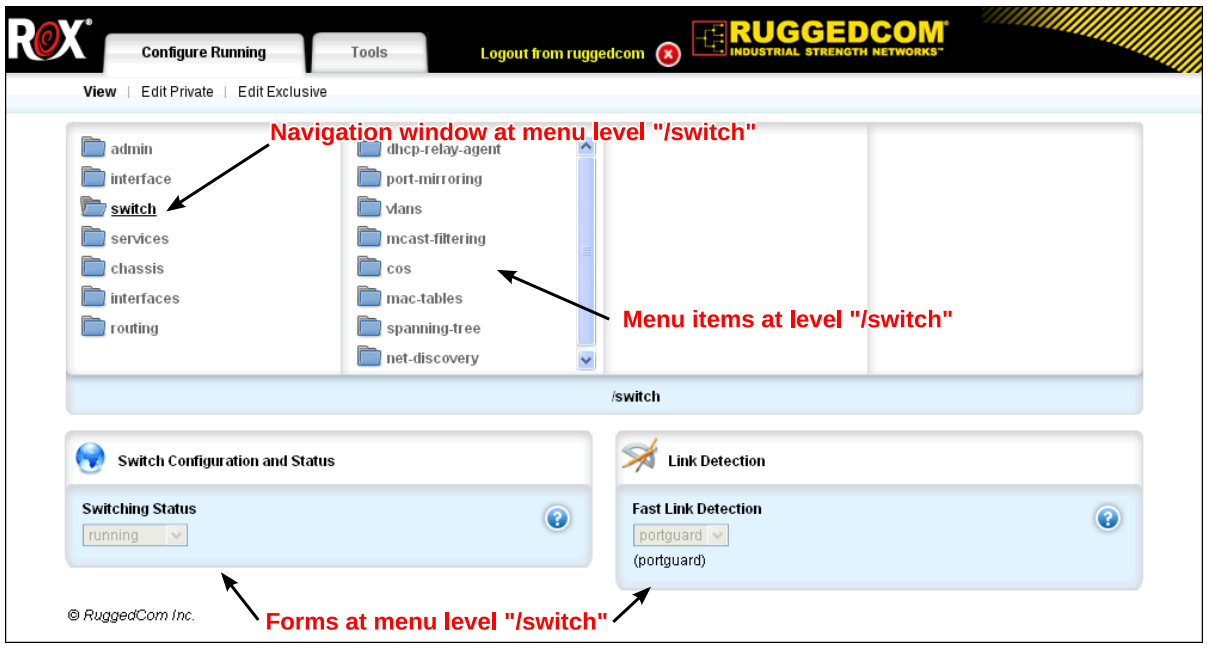

Figure 1.2. The ROX™ Web Interface

Menu items listed in a panel of the navigation window at a given point in the menu hierarchy may be:

- Submenus, which are marked using the icon, or
- Actions, which are marked using the icon.

Note that green submenu icons represent operational data.

Tables and forms relevant to the selected menu item appear below the navigation window.

The icons in the upper left corner of the forms and tables are used to signify the type of content represented in each form or table.

- The green arrow icon signifies operational data.
- The key icon signifies the key in key settings.
- The blue globe icon signifies the global group (a high-level grouping of items).
- The pencils and protractor icon signifies configuration data.
- The paper and pencil icon signifies results. This icon is usually found on a form where there are parameters to enter.

Every web page in the ROX™ user interface has a header, illustrated above, containing:

- The ROX™ and RuggedCom logos and a **Logout** button, which terminates the current web session.
- The tabs: *Configure Running* and *Tools*.

The *Configure Running* tab selects the configuration interface described above. A menu bar below the page header displays the following editing mode controls:

- *View*: View configuration settings only.
- *Edit Private*: Enter a configuration editing mode where you can make changes to the system. Your changes are applied to the active system only when you commit them. Edit sessions are self contained: the changes made in your edit session are not visible to other users in other edit sessions.
- *Edit Exclusive*: Enter a configuration editing mode where, after committing your changes, you can specify a timeout period to test the changes. At the end of the timeout period, your changes to revert back to the original settings. Use this mode when you want to test changes before committing them permanently. When you click **Commit**, a dialog prompts you to set a commit timeout. Type a value and select a unit of time. ROX temporarily applies your changes to the active system for the specified time. To cancel the commit and discard the changes, click **Abort Commit** before the time elapses. To permanently commit your changes, click **Commit** before the time elapses.

In many cases, the tables appear on a screen closer to the top level and clicking on one of the submenus brings up the form(s) associated with the table. For example, clicking on the Chassis menu and then the Hardware submenu will display the Slot Hardware table. Further clicking on the pm1 submenu will display the Slot Hardware form.

The *Tools* tab displays a menu of tools in the menu bar, with the following structure:

- Device Info: displays text from various system logs. You can specify the number of lines to view and a text filter.
  - Messages Viewer: displays all events from **`/var/log/messages`**.
  - Syslog Viewer: displays syslog events from **`/var/log/syslog`**.
  - Authlog Viewer: displays authentication events from **`/var/log/auth.log`**.
  - Layer2log Viewer: displays Layer 2 events from **`/var/log/layer2`**.
  - Kernlog Viewer: displays kernel events from **`/var/log/messages`**.
- Accessories
  - Ping: an ICMP echo tool for IPv4 addresses
  - Ping6: an ICMP echo tool for IPv6 addresses
  - Tcpdump: a packet analyzer for TCP/IP and other packets

  - Traceroute: a tool for displaying route or path information and packet transit delays between IPv4 addresses
  - Traceroute6: a tool for displaying route or path information and packet transit delays between IPv6 addresses
- CLI: a command line interface window
- Users: displays a list of currently connected users, provides controls to kick users off of the system, and provides a message board to send messages to users.
- Upload: uploads configuration files, feature keys, elan certificates, ipsec certificates, ca certificates, and crl certificates to the system from your workstation. From the *Choose file type* list, select the type of file to upload. Click **Choose File** and navigate to and select a file on your workstation. To upload the selected file, click **Send**.
- Download: downloads configuration files, feature keys, elan certificates, ipsec certificates, ca certificates, crl certificates, log files, and rollback files from the system to your workstation. From the *Choose file type* list, select the type of file to download. Click **List files**; a list of available files appears. To download a file, right-click on a file name and select **Save Link As** (the name of the menu option will vary, depending on your browser). To open a file in a new window or tab, click on a file name.

1.2.1. Top-level Menu Categories

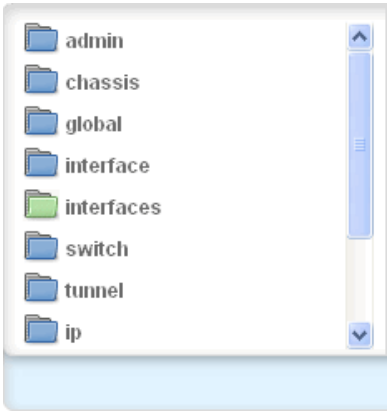

Figure 1.3. Top-level Menu

Below is a description of the categories in the top-level menu that is shown above.

admin

The admin menu is used for configuring functions related to the administration of the router. Functions include DNS, alarms, logging, authentication, users, software upgrade, notifications and SNMP.

chassis

The chassis menu is used for configuring the chassis.

global

The global menu is used for configuring global functions including profiles for PPP and cellular modems.

interface

The interface menu is used for configuring the interface, including (where applicable) sections for WAN, serial, modem and trunks.

interfaces

The interfaces menu displays the status of functions configured via the interface menu. For example, eth functions can be configured using the eth submenu that is accessible from the interface menu. The eth status can be viewed by clicking on the eth submenu of the interfaces menu.

switch

The switch menu is used for configuring Layer 2 packet switching functions. Functions included are port security, DHCP relay agent, port mirroring, multicast filtering, CoS, mac tables, spanning tree, VLANs, layer 3 switching and net discovery. You can also reset switched ports and clear switched port statistics and cable diagnostic test results.

tunnel

The tunnel menu is used for configuring IP tunnels using IPsec, Layer 2 tunnelling functions and Generic Routing Encapsulation (GRE).

ip

The ip menu is used for configuring the ROX™ system's IP network interfaces.

qos

The qos menu is used for configuring traffic control.

routing

The routing menu is used for configuring the routing features. Included are sections on dynamic, static, status and multicast routing.

security

The security menu is used for configuring security, including the firewall.

services

The services menu is used for configuring various services. These services include timekeeping, VRRP, DHCP server and linkfailover.

1.3. Making Configuration Changes

To make configuration changes, select *Edit Private* or *Edit Exclusive* mode from the configuration view. The same navigation window, tables and forms are redisplayed, but with additional controls, as illustrated below.

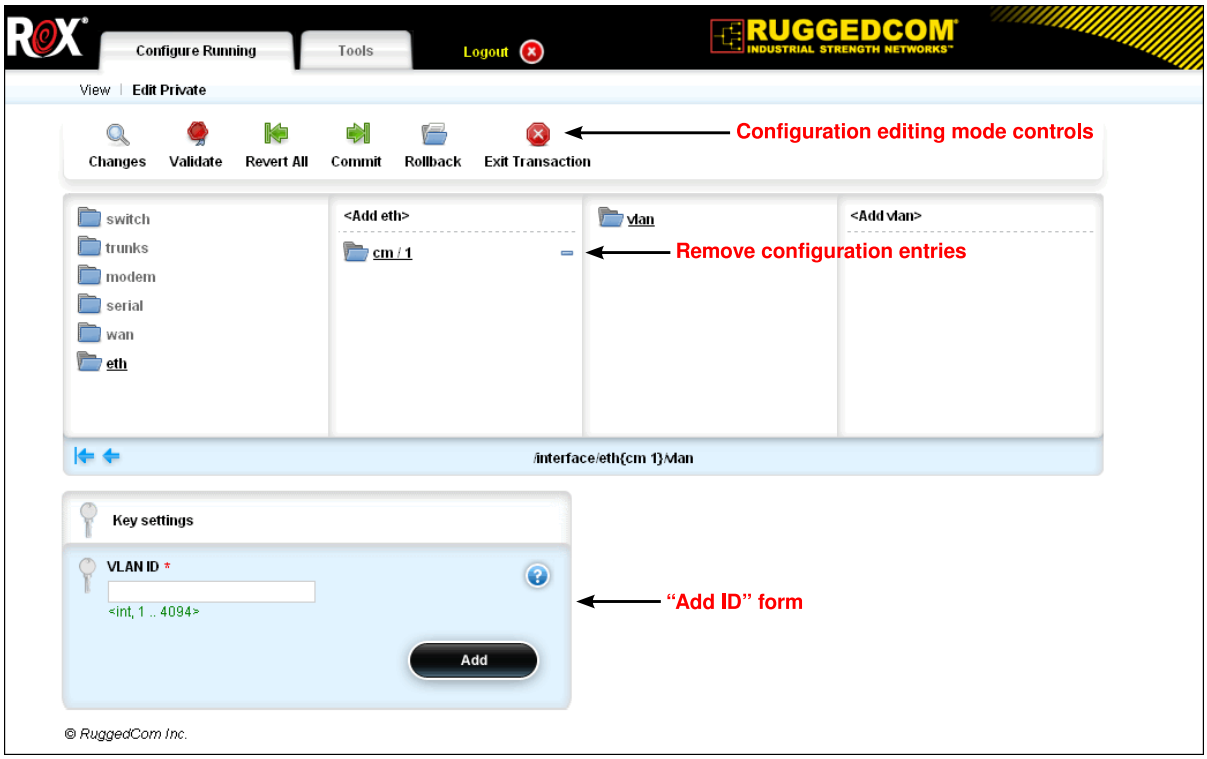

Figure 1.4. Example of Edit Private Mode

The example above depicts the process of adding a VLAN ID to an interface. The **interface/eth/cm1** menu can be seen to contain:

- A configuration entry, followed by a "delete" icon, , which removes the corresponding entry.

Clicking on <add vlan> displays the *Add ID* form below the navigation window, which prompts for a VLAN ID. Entering a VLAN ID and clicking **Add** adds the selected VLAN to the currently selected interface.

Note the help button, , on the *Add ID* form which, when clicked, displays context-sensitive information about the corresponding data field.

In Edit Private and Edit Exclusive mode, a red asterisk appears beside fields that are mandatory for configuration. Note the red asterisk next to the field name (VLAN ID) in the Key settings form.

Several controls below the header and menu bar are used to affect the behaviour of the changes made during the current configuration editing session:

Changes

Present a summary of all pending changes.

Validate

Automatically check the validity of pending changes.

Revert All

Abort all pending changes.

Commit

Commit all pending changes - save changes persistent configuration storage and to the running system.

Rollback

Present a list of change sets made to date, with an option to revert a selected set of changes.

Exit Transaction

Exit from configuration editing mode. If there are pending changes, a prompt will be presented to verify the discarding of all pending changes.

1.3.1. Configuring Tables Using Key Settings Forms

Much of the information in ROX™ is organized into tables. Each table is indexed or sorted by a key, which is a piece of information such as a name, address, or other variable. For example, a Chassis Hardware table is indexed by slot name (with the slot name being the key) and a DNS Server table is indexed by IP address (with the IP address being the key). Key information can be added using the key settings forms. To add server information to a DNS server table, for example, add the server address to the key settings form and this information will appear in the DNS server table.

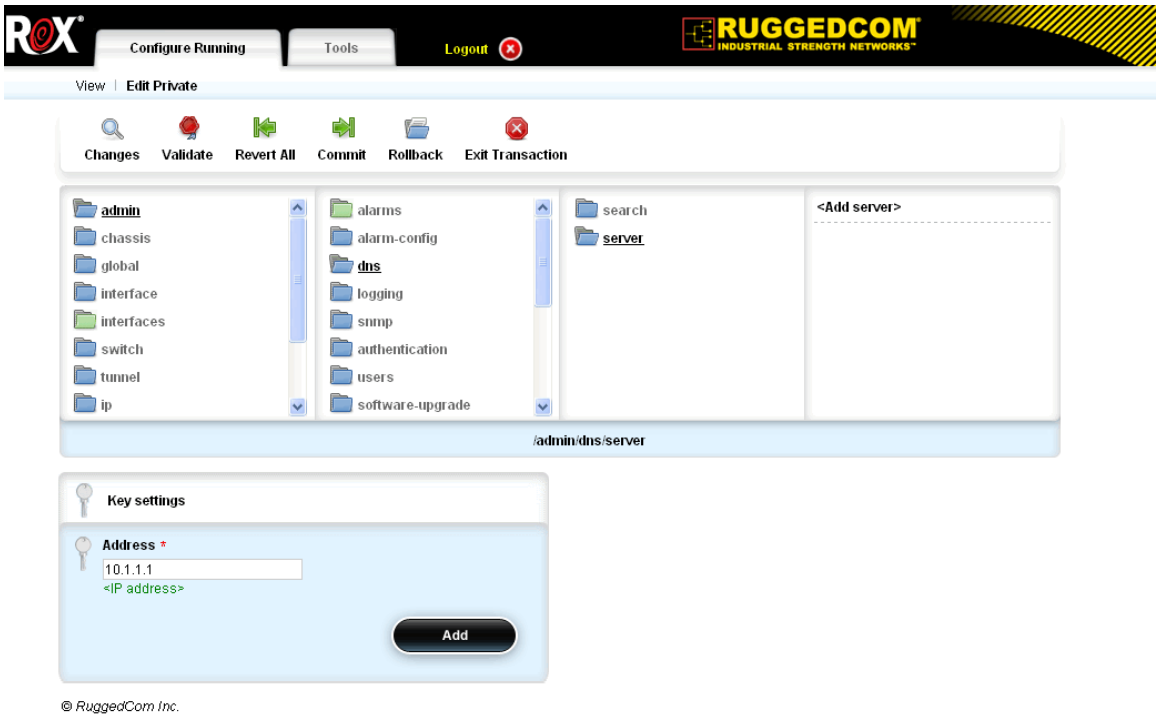

Figure 1.5. Adding Key Information

To add key information to a table, enable *Edit Private* or *Edit Exclusive* mode, enter the information on the key settings form, and click *Commit*. To return to the *View* mode after committing your changes, click *Exit Transaction*.

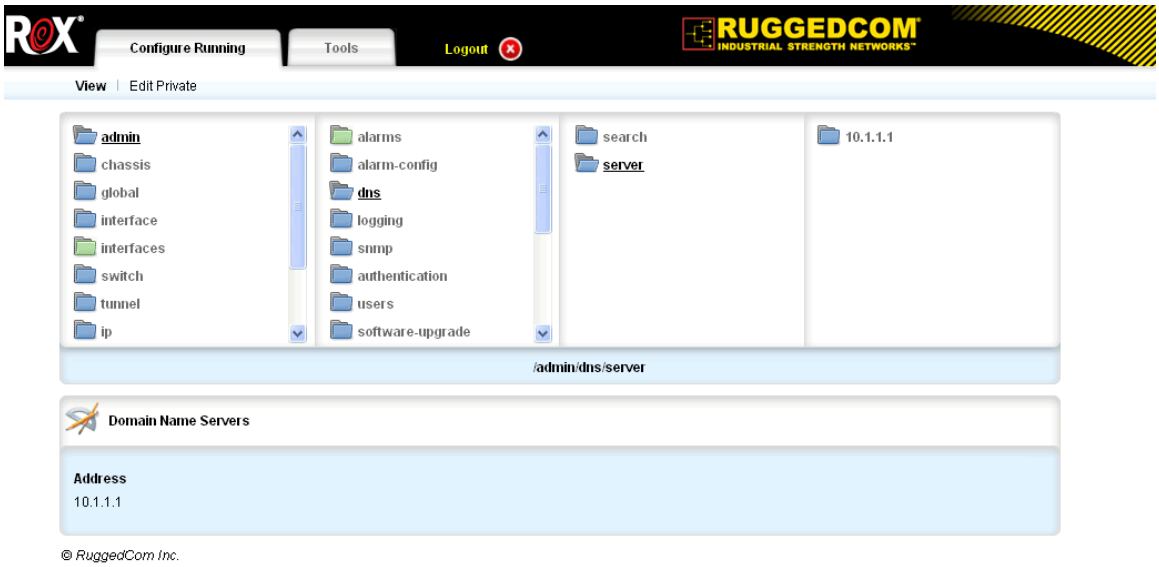

Figure 1.6. Key Information in a Table

The information entered in the key settings form will now appear in the table. Note that the table appears on the server screen, while the key settings form appears on the address screen, which is a submenu linked to the server screen (see below).

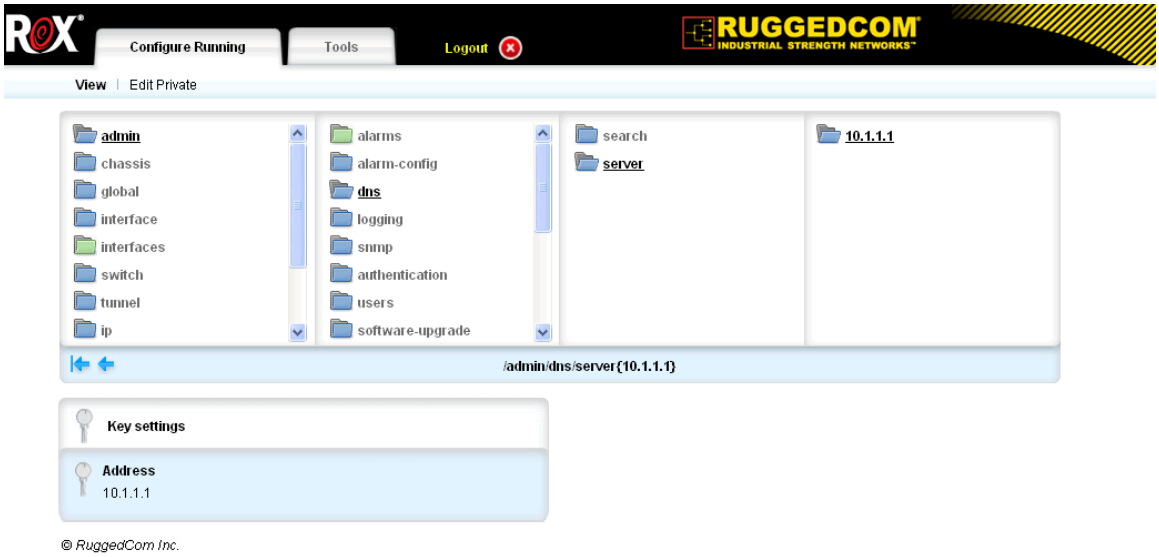

Figure 1.7. Example of Key Settings 1

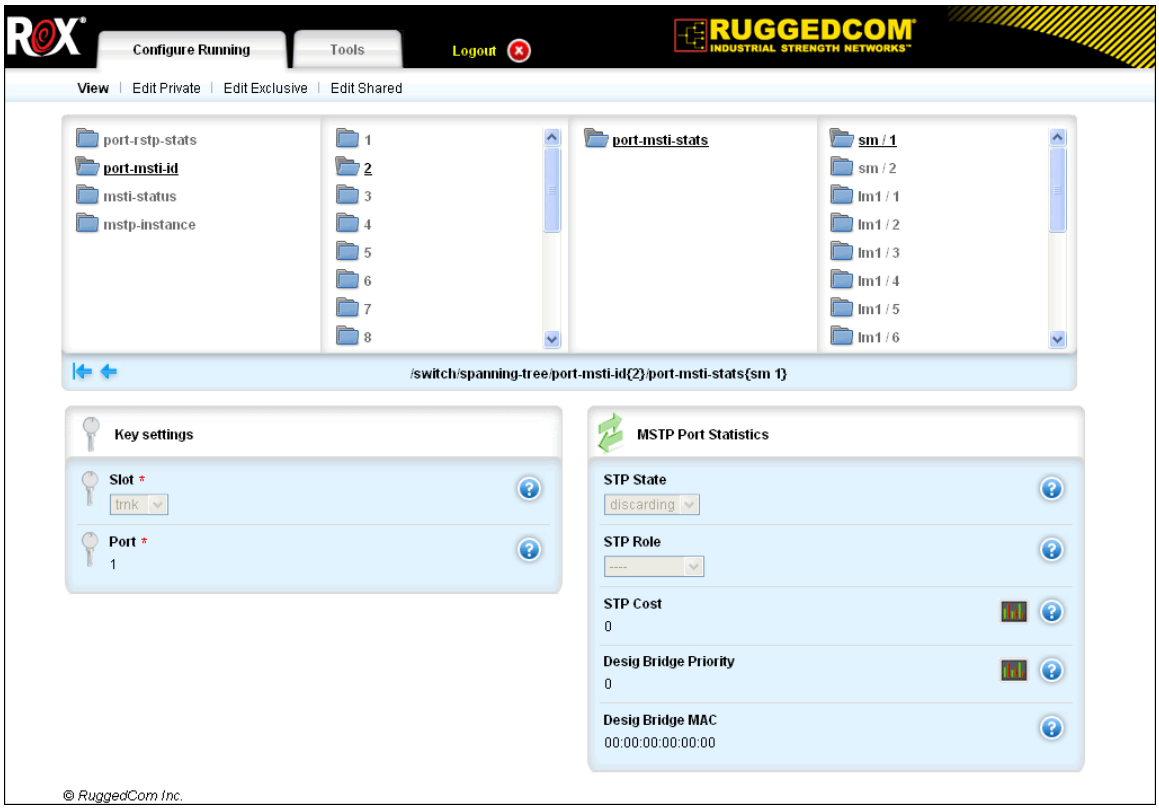

Figure 1.8. Example of Key Settings 2

The submenus that display the key settings forms appear in the far right column of the screen. Sometimes, it will be necessary to traverse several menu screens to get to a key settings form.

1.3.2. Viewing More Information in Tables

Occasionally, a table may have more entries that are not visible in the initial view. If you encounter a table that has a line of linked text at the top with the word "Next", and a number in parentheses ( ), you can click on the "Next" link and access additional entries. The two figures below illustrate this situation. In this case, there are 18 entries in the table. The first table contains 16 entries and 2 entries follow in the next table.

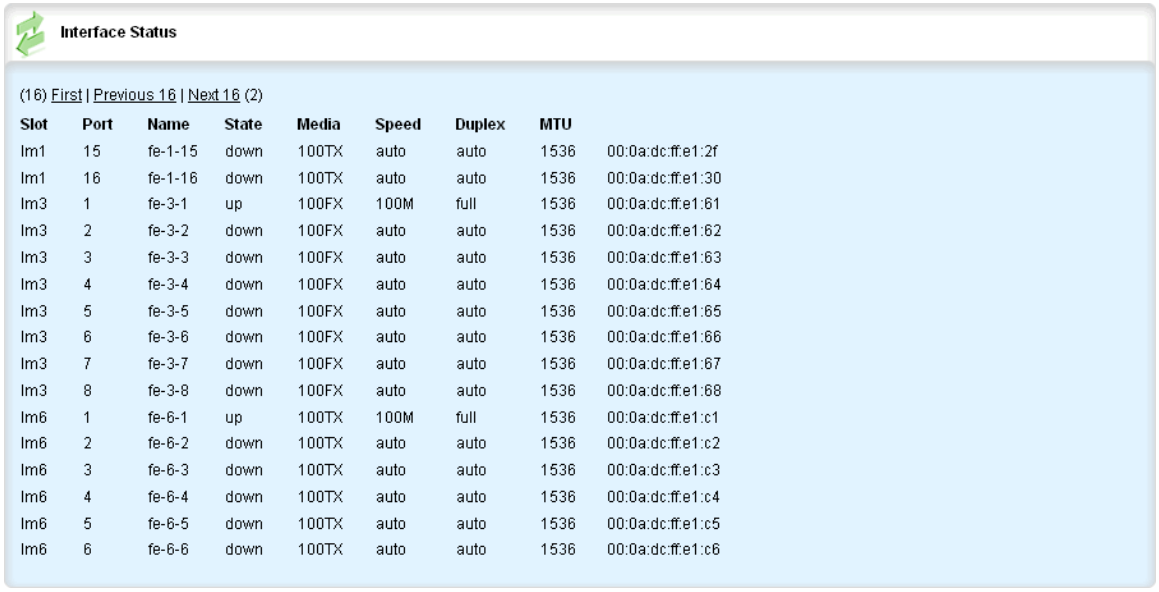

Interface Status

(16) First | Previous 16 | Next 16 (2)

Slot	Port	Name	State	Media	Speed	Duplex	MTU	
lm1	15	fe-1-15	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:2f
lm1	16	fe-1-16	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:30
lm3	1	fe-3-1	up	100FX	100M	full	1536	00:0a:dc:ff:e1:61
lm3	2	fe-3-2	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:62
lm3	3	fe-3-3	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:63
lm3	4	fe-3-4	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:64
lm3	5	fe-3-5	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:65
lm3	6	fe-3-6	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:66
lm3	7	fe-3-7	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:67
lm3	8	fe-3-8	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:68
lm6	1	fe-6-1	up	100TX	100M	full	1536	00:0a:dc:ff:e1:c1
lm6	2	fe-6-2	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c2
lm6	3	fe-6-3	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c3
lm6	4	fe-6-4	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c4
lm6	5	fe-6-5	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c5
lm6	6	fe-6-6	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c6

Figure 1.9. First Table of Information

Interface Status

(18) First | Previous 16

Slot	Port	Name	State	Media	Speed	Duplex	MTU	
lm6	7	fe-6-7	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c7
lm6	8	fe-6-8	up	100TX	100M	full	1536	00:0a:dc:ff:e1:c8

Figure 1.10. Second Table of Information

The second table of information shows the balance of the entries and contains a link back to the previous entries.

2. System Administration

This chapter describes administration-related functions and the Administration menu. Information on the Administration submenus is found throughout Part 1 of this guide.

2.1. Administration menu

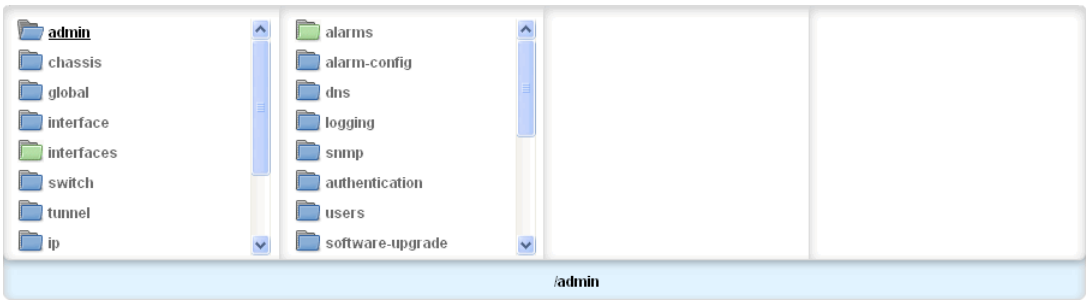

Figure 2.1. Administration menu

The Administration (Admin) menu is accessible from the main menu. Use this menu to link to submenus related to alarms, DNS, logging, SNMP, authentication, user IDs and passwords, software versions (upgraded) and netconf.

As well, you can link directly from the Admin menu to commands called "actions" (see below) that enable you to perform various functions, including the following: clearing all alarms; acknowledging all alarms; shutting down the system; rebooting the system; setting the system clock; and restoring the factory defaults.

2.2. System Commands

This section describes where to find basic system commands using the Administration menu and its menu actions. The following forms are accessible from the Administration menu.

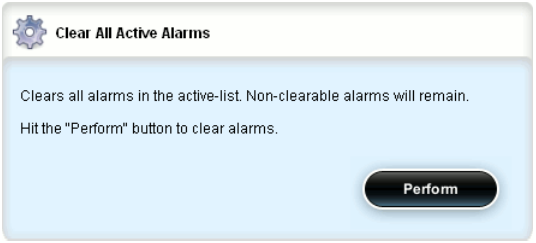
Clear All Active Alarms
Clears all alarms in the active-list. Non-clearable alarms will remain.
Hit the "Perform" button to clear alarms.
Perform

Figure 2.2. Clear All Alarms Menu Action form

To clear all alarms, click on the *clear-all-alarms* menu action and then click the *Perform* button on the Clear All Alarms form.

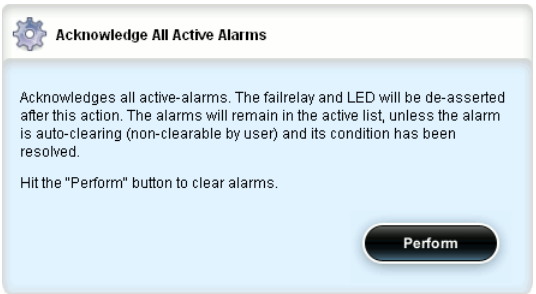
Acknowledge All Active Alarms
Acknowledges all active-alarms. The failrelay and LED will be de-asserted after this action. The alarms will remain in the active list, unless the alarm is auto-clearing (non-clearable by user) and its condition has been resolved.
Hit the "Perform" button to clear alarms.
Perform

Figure 2.3. Acknowledge All Alarms Menu Action form

To acknowledge all alarms, click on the *acknowledge-all-alarms* menu action and then click the *Perform* button on the Acknowledge All Alarms form.

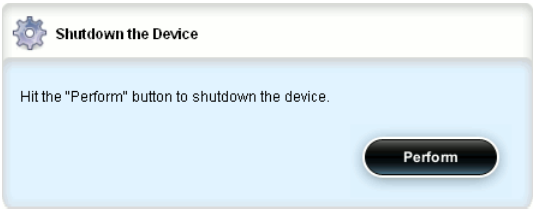

Figure 2.4. Shutdown the Device Menu Action form

To shut down the device, click on the *shutdown* menu action and then click the *Perform* button on the Shutdown the Device form.

The device never enters a permanent shutdown state. If you instruct the device to shutdown, it shuts down and provides a time-out period during which you can remove power from the device. The default time-out period is 300 seconds (five minutes). At the end of the time-out period, the device reboots and restarts.

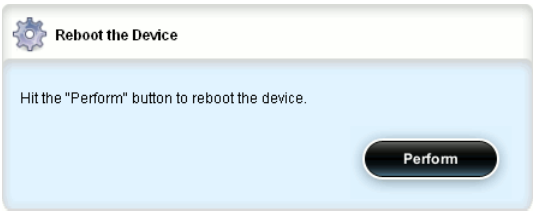

Figure 2.5. Reboot the Device Menu Action form

To reboot the device, click on the *reboot* menu action and then click the *Perform* button on the Reboot the Device form.

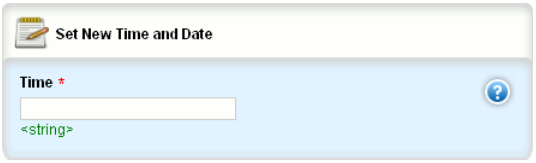

Figure 2.6. Set New Time and Date form

The Set New Time and Date form configures the current time and date settings.

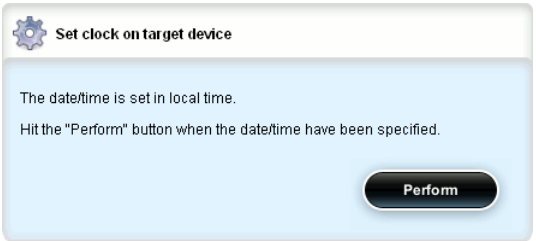
Set clock on target device
The date/time is set in local time.
Hit the "Perform" button when the date/time have been specified.
Perform

Figure 2.7. Set Clock on Target Device form

To set the clock on the target device, click on the *setSystemClock* menu action, then enter the relevant time/date information into the Set New Time and Date form. The information must be in the following format: YYYY-MM-DD HH:MM:SS. After entering this information, click the *Perform* button on the Set clock on target device form.

For more detailed information on time synchronization, refer to Chapter 3, *Time Synchronization*.

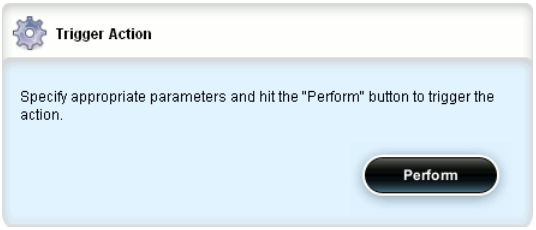

Figure 2.8. Restore-factory-defaults Trigger Action form

To restore factory defaults to the system, click on the *restore-factory-defaults* menu action and then click the *Perform* button on the Restore-factory-defaults Trigger Action form.

The Administration, Hostname, Timezone and Current System Time forms are accessible from the Admin menu.

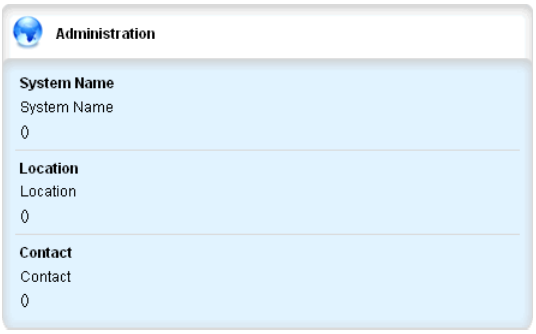

Figure 2.9. Administration form

System Name

Synopsis: *A string*

Default: *System Name*

An administratively-assigned name for this managed node. By convention, this is the node's fully-qualified domain name. If the name is unknown, the value is the zero-length string.

Location

Synopsis: *A string*

Default: *Location*

The physical location of this node (e.g., 'telephone closet, 3rd floor'). If the location is unknown, the value is the zero-length string.

contact

Synopsis: *A string*

Default: *Contact*

The textual identification of the contact person for this managed node, together with information on how to contact this person. If no contact information is known, the value is the zero-length string.

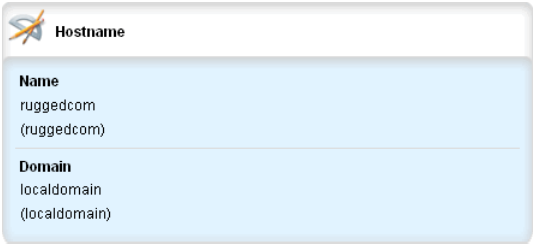

Figure 2.10. Hostname form

The hostname is the name of the product. (This can be changed, though.)

name

Synopsis: *A string conforming to: "[A-Za-z0-9]([A-Za-z0-9-]*[A-Za-z0-9])*"*

Default: *ruggedcom*

The hostname is the name of this device.

domain

Synopsis: *Domain name (RFC 1034)*

Default: *localdomain*

The domain for this hostname.

Figure 2.11. Timezone form

Timezone Category

Synopsis: *string*

Selects the timezone. Note that the Etc/GMT timezones conform to the POSIX style and have their signs reversed from common usage. In POSIX style, zones west of GMT have a positive sign; zones east of GMT have a negative sign.

Timezone

Synopsis: *string*

Selects the timezone.

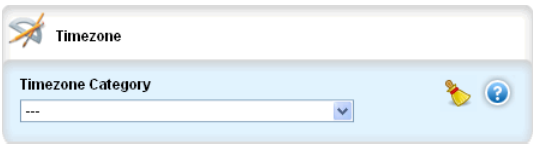

Figure 2.12. Setting the Timezone Form - in Edit Private Mode

To set the time zone, enter Edit Private mode and click on the Timezone Category field. Use the drop-down menu which appears to select the appropriate time zone. Daylight saving time will adjust automatically, if applicable to your zone.

Figure 2.13. Current System Time form

The Current System Time form displays the current time.

UTC Time

Synopsis: *string*

The current GM Time

Local Time

Synopsis: *string*

The current local time

2.3. Administrative Access Control

The following access control forms are accessible from the Administration menu - by clicking on the main menu under *admin*.

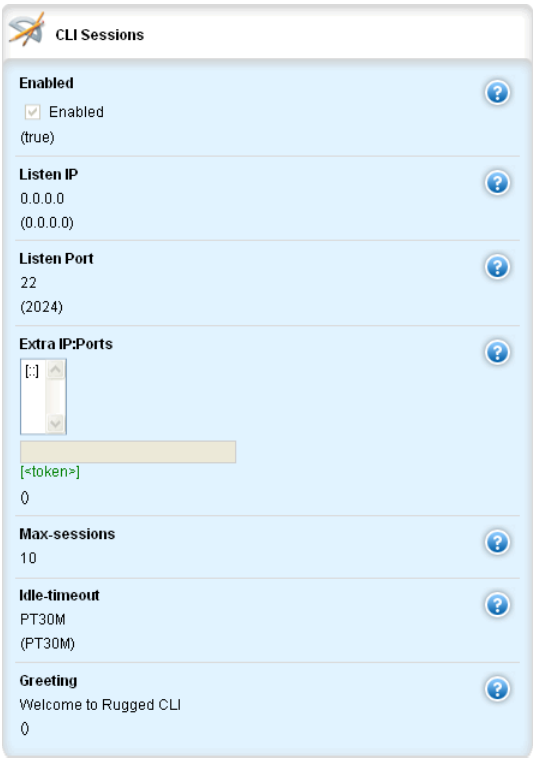

Figure 2.14. CLI Sessions form

enabled

Synopsis: *boolean*

Default: *true*

Provides the ability to configure CLI features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the CLI will listen on for CLI requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *22*

The port on which the CLI listens for CLI requests. The default is port 22.

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The CLI will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of CLI Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *10*

The maximum number of concurrent CLI sessions

Idle Timeout

Default: *PT30M*

Maximum idle time before terminating a NETCONF session. If the session is waiting for notifications, or has a pending confirmed commit, the idle timeout is not used. The default value is 0, which means no timeout.

greeting

Synopsis: *string*

Default: *Welcome to Rugged CLI*

Sets the greeting presented when the user logs in to the CLI.

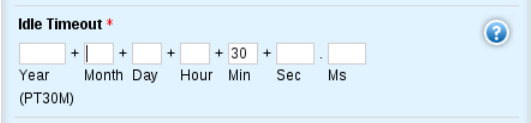

Figure 2.15. Idle-timeout field

Clicking on the Idle-timeout field on the CLI Sessions form allows you to choose a value for this field. The default value is PT30M, which stands for "Precision Time 30 Minutes". This refers to the time when an inactive session expires or times out. Only integer values corresponding to the following fields can be entered: Year, Month, Day, Hour, Min, Sec, or Ms. The example above shows the default value of PT30M, which corresponds to the Min field.

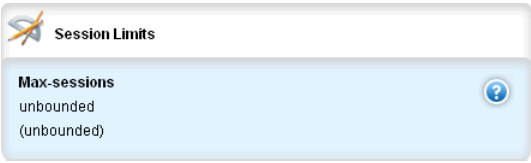

Figure 2.16. Session Limits form

The Session Limits form is used for setting the maximum number of users sessions on a northbound channel.

Maximum Sessions Total

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *70*

Puts a limit to the total number of concurrent sessions to ROX 2.

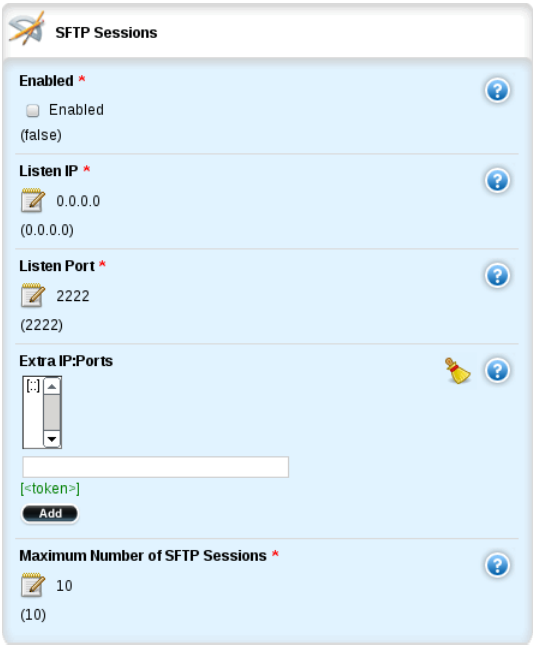

Figure 2.17. SFTP Sessions form

The SFTP Sessions form sets the parameters for Secure File Transfer Protocol (SFTP) sessions.

enabled

Synopsis: *boolean*

Default: *false*

Enable/Disable the SFTP user interface

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the SFTP will listen on for SFTP requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *2222*

The port the SFTP will listen on for SFTP requests (default 2222).

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The SFTP will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of SFTP Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *10*

The maximum number of concurrent SFTP sessions

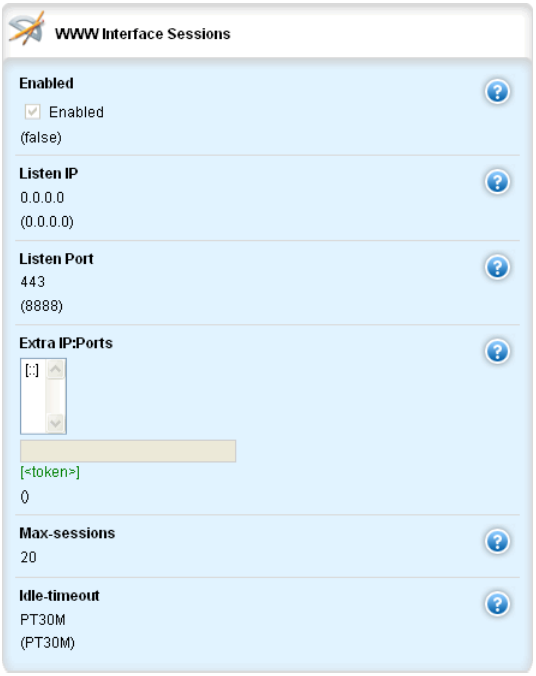

Figure 2.18. WWW Interface Sessions

The WWW Interface Sessions form provides control of WWW User Interface settings.

enabled

Synopsis: *boolean*

Default: *true*

Provides the ability to configure WebUI features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the CLI will listen on for WebUI requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *443*

The port on which the WebUI listens for WebUI requests. The default is port 443.

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The WebUI will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of WebUI Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *20*

The maximum number of concurrent WebUI sessions

Idle Timeout

Default: *PT30M*

Maximum idle time before terminating a WebUI session. If the session is waiting for notifications, or has a pending confirmed commit, the idle timeout is not used. The default value is 0, which means no timeout.

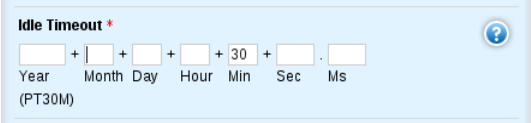

Figure 2.19. Idle-timeout field

Clicking on the Idle-timeout field on the WWW Interface Sessions form allows you to choose a value for this field. The default value is PT30M, which stands for "Precision Time 30 Minutes". This refers to the time when an inactive session expires or times out. Only integer values corresponding to the following fields can be entered: Year, Month, Day, Hour, Min, Sec, or Ms. The example above shows the default value of PT30M, which corresponds to the Min field.

2.4. User Accounts

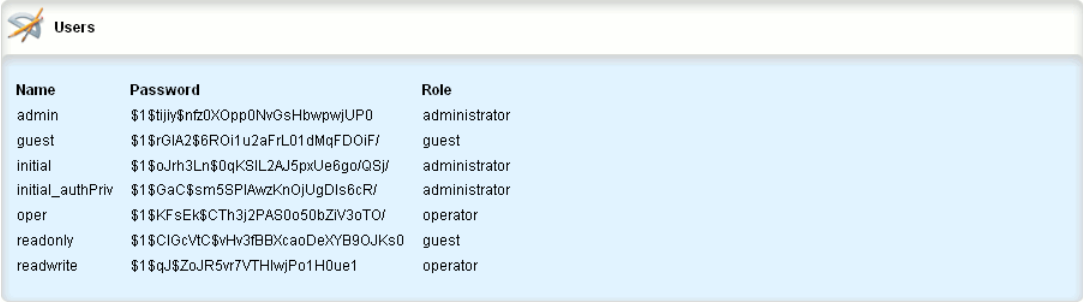

Figure 2.20. Users menu

The Users menu is accessible from the main menu under *admin*. This menu is used to access commands needed for creating and managing passwords for administrators, operators and guests. Both private and public passwords can be created. The Admin Users ID Table (below) can be found on the same screen as the Users menu. Clicking on admin, guest, oper, private or public will lead you to the Users ID forms for each of these options.

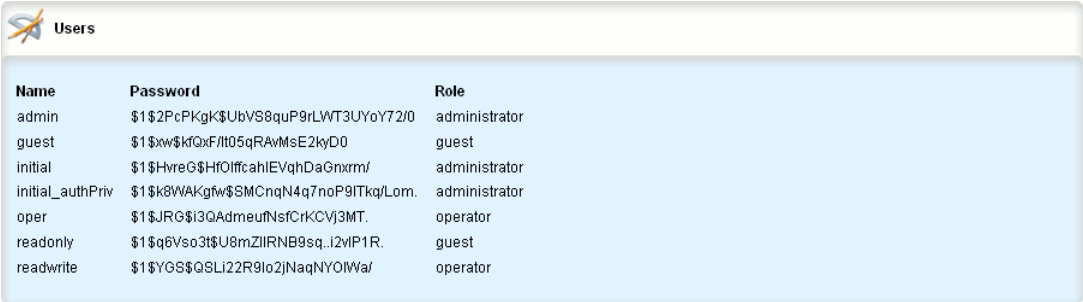

Figure 2.21. Users table

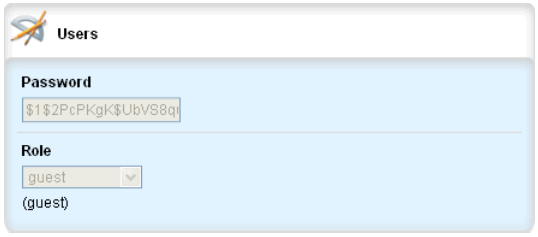

Figure 2.22. Users form

name

Synopsis: *string*

User Name

password

Synopsis: *A string*

User Password

role

Synopsis: *string - one of the following keywords { guest, operator, administrator }*

Default: *guest*

User Role

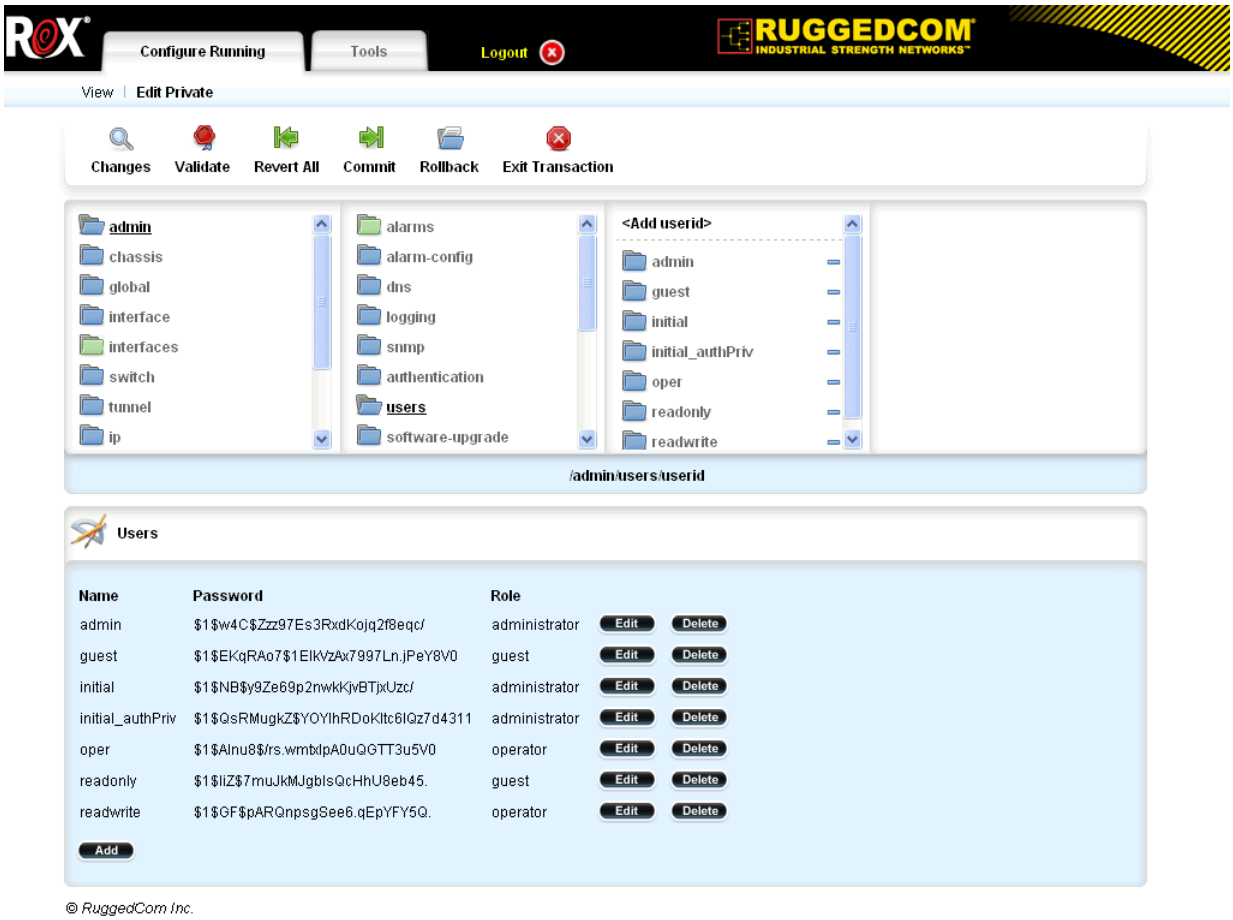

Figure 2.23. Users Screen in Edit Private View

Passwords can be managed, added and deleted while in the Edit mode.

2.5. Software Upgrade

ROX™ supports two system partitions. One is always active and the other is inactive. ROX™ always applies software upgrades to the inactive partition, providing the following advantages:

1. The current system is unaffected and can operate normally while the upgrade is in progress
2. The current partition remains intact, allowing you to roll back to the original system if needed

After a successful upgrade, the next reboot boots the upgraded partition.

The following applies to software upgrades:

- All system configurations and all user files (featurekeys, configuration files etc.) are carried over to the upgraded partition.
- All configurations are locked during an upgrade and until the upgraded partition is booted. This prevents post-upgrade configuration changes that are not carried over to the upgraded partition.
- Completed upgrades can be declined before the next reboot.
- If major system failures are detected upon booting the upgraded partition, the system will automatically roll back to the previous partition.

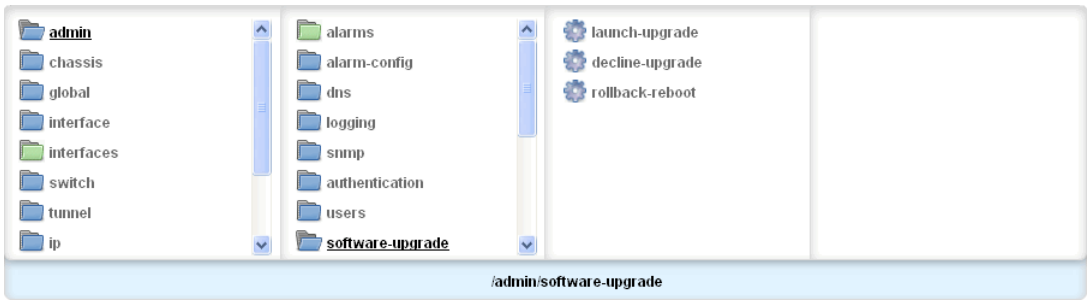

Figure 2.24. Software-Upgrade menu

The Software-Upgrade menu is accessible from the main menu under *admin*. The path to this menu is *admin/software-upgrade*. This menu links to functions that will enable the user to upgrade software, launch the upgraded software, decline new upgrades, and rollback and reboot. The Upgrade Monitoring form and Upgrade Settings form appear on the same screen as the Software-Upgrade menu.

Figure 2.25. Upgrade Settings

In edit mode, define an upgrade server on the *Upgrade Settings* form by setting the Server URL and Target ROX Version parameters. The Upgrade Server URL is the location of the ROX™ software repository. Target ROX Version is the version of ROX to which you are upgrading. For information on setting up an upgrade server, see Appendix C, *Setting Up An Upgrade Server*.

Upgrade Server URL

Synopsis: *string*

repository-url

Target ROX Version

Synopsis: *string*

target-version

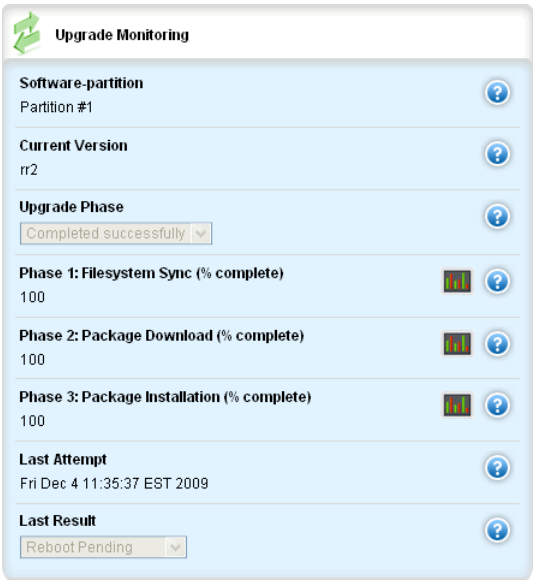

Figure 2.26. Upgrade Monitoring

The Upgrade Monitoring form displays the status of the current upgrade operation.

software-partition

Synopsis: *A string*

The current active partition number. The unit has two software partitions: #1 and #2. Upgrades are always peformed to the other partition.

Current Version

Synopsis: *A string*

The current operating software version.

Upgrade Phase

Synopsis: *string - one of the following keywords { Failed, Completed successfully, Unknown state, Installing packages, Downloading packages, Copying filesystem, Estimating upgrade size, Inactive }*

The current phase or state of the upgrade. It is one of 'Estimating upgrade size', 'Copying filesystem', 'Downloading packages', 'Installing packages', Unknown state', 'Completed successfully', or 'Failed'. These phrases will not vary, any may be used programmitcally for ascertaining state.

status-message

Synopsis: *string*

Additional details on the status of the upgrade

Phase 1: Filesystem Sync (% complete)

Synopsis: *integer*

Phase 1 of the upgrade involves synchronizing the filesystem with the partition to which we are upgrading.

This reflects the estimated percent complete.

Phase 2: Package Download (% complete)

Synopsis: *integer*

Phase 2 of the upgrade downloads all packages that require an update. This reflects the estimated percent complete.

Phase 3: Package Installation (% complete)

Synopsis: *integer*

Phase 3 of the upgrade installs all packages that require an update. This reflects the estimated percent complete.

Last Attempt

Synopsis: *A string*

The date and time of completion of the last upgrade attempt.

Last Result

Synopsis: *string - one of the following keywords { Interrupted, Declined, Not Applicable, Reboot Pending, Unknown, Upgrade Failed, Upgrade Successful }*

Indicates whether or not the last upgrade completed successfully

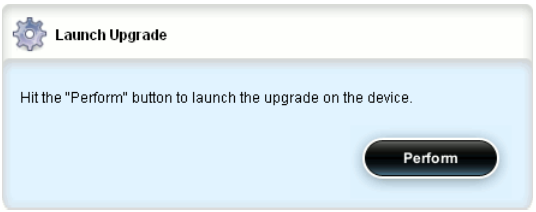

Figure 2.27. Launch Upgrade

To launch an upgrade, click on the *launch-upgrade* menu action and then click the *Perform* button on the Launch Upgrade form. Note that the server URL and version name information must be entered in the Upgrade Settings form prior to launching the upgrade. For detailed step-by-step instructions on how to perform a software upgrade, refer to Appendix A, *Upgrading Software*.

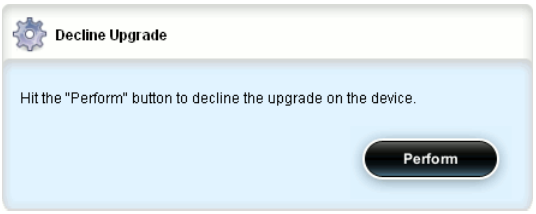

Figure 2.28. Decline Upgrade

To decline an upgrade, click on the *decline-upgrade* menu action and then click the *Perform* button on the Decline Upgrade form.

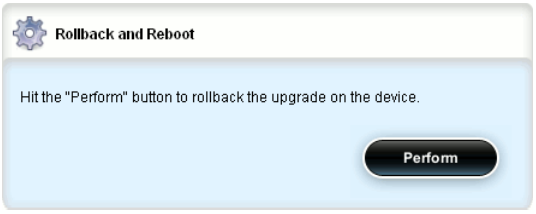

Figure 2.29. Rollback and Reboot

To roll back an upgrade, click on the *rollback-reboot* menu action and then click the *Perform* button on the Rollback and Reboot form.

Rollback and Reboot “rolls back” the system to the previously active software installation, which is stored on the alternate of two filesystem partitions in flash memory. Performing this action will result in rebooting the system using the old software installation along with its configuration.

> *Any configuration changes made since the last software upgrade will not be reflected after rebooting to the "rolled-back" software installation.*

2.6. ROXflash Cross-Partition Imaging Tool - Software Downgrade

ROX™ supports two system partitions. One is always active and the other is inactive. ROXflash allows you to flash any ROX™ software version to the inactive partition.

To obtain a flash image, contact your RuggedCom sales representative. Place the flash image in a location on your network accessible to the ROX™. On the ROXflash form, enter the URL for the flash image and flash it to the inactive partition. The flash image will be active after the next reboot.

2.6.1. Uses

Use ROXflash for downgrading to an earlier version of the ROX software. For example, your organization has certified a specific version of the ROX software, and all ROX™ units must run the certified version. Due to an equipment issue, you need to install a new ROX™ unit that comes with a later version of the software. In this example, use ROXflash to install the earlier version of the software on the new unit.

Use ROXflash only to install earlier versions of the ROX software. Software upgrades to later versions should be performed using the Software Upgrade function.

Table 2.1, “Differences Between ROXflash and Software Upgrade Functions” outlines some of the key differences between the ROXflash and Software Upgrade functions. For more information on the Software Upgrade function, see Section 2.5, “Software Upgrade”.

ROXflash	Software Upgrade
Used primarily for downgrades.	Used only for upgrades; does not support downgrades (except for rollbacks).
Uses a flash image ordered from a RuggedCom Sales Representative.	Uses an archive of ROX™ software packages hosted on an upgrade server. The archive is available on RuggedCom.com for download.
Downgrades to any software version supplied in an image.	Rolls back only to the last version stored on the alternate partition.
Does not transfer system configurations and files to the next software version. ROXflash returns the unit to its factory default settings. Configurations must be reloaded after rebooting.	Transfers configurations and files to the upgraded software version; reverts to the previous configurations in a rolled back version.

Table 2.1. Differences Between ROXflash and Software Upgrade Functions

2.6.2. ROXflash Configuration

Figure 2.30. ROX-Imaging menu

The ROX-Imaging menu is accessible from the main menu under *admin*. The ROXflash Monitoring form appears on the same screen as this menu.

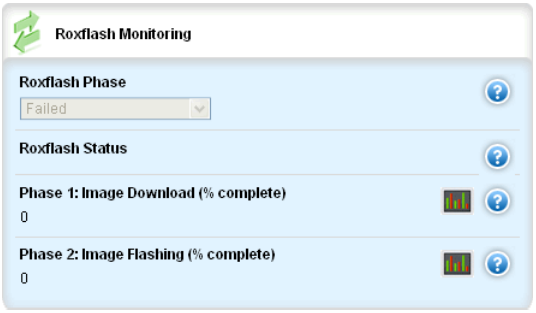

Figure 2.31. ROXflash Monitoring form

This form shows the progress and state of the roxflash operation (during an upgrade or downgrade).

ROXflash Phase

Synopsis: *string - one of the following keywords { Failed, Completed successfully, Unknown state, Imaging partition, Downloading image, Inactive }*

The current phase or state of the ROXflash operation. It is always one of: 'Inactive', 'Downloading image', 'Imaging partition', 'Unknown state', Completed successfully, or 'Failed'. These phrases do not vary, and may be used programatically for ascertaining state.

ROXflash Status

Synopsis: *A string*

Detailed messages about ROXflash progress.

Phase 1: Image Download (% complete)

Synopsis: *integer*

Phase 1 of ROXflash downloads the image from a URL. This reflects percent complete.

Phase 2: Image Flashing (% complete)

Synopsis: *integer*

Phase 2 of ROXflash flashes the image to the alternate partition. This reflects percent complete.

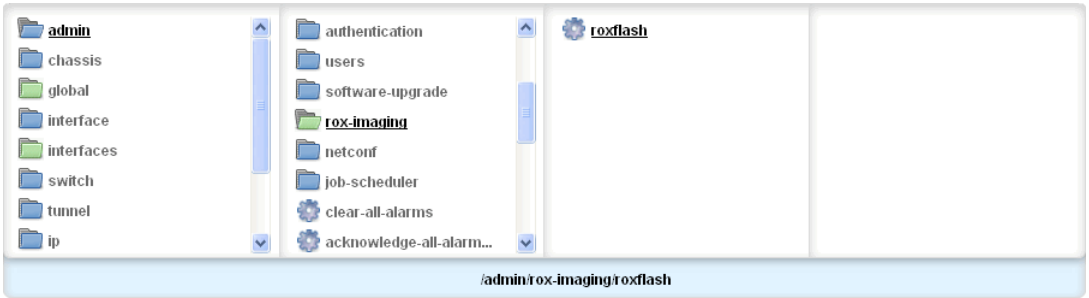

Figure 2.32. ROXFlash menu

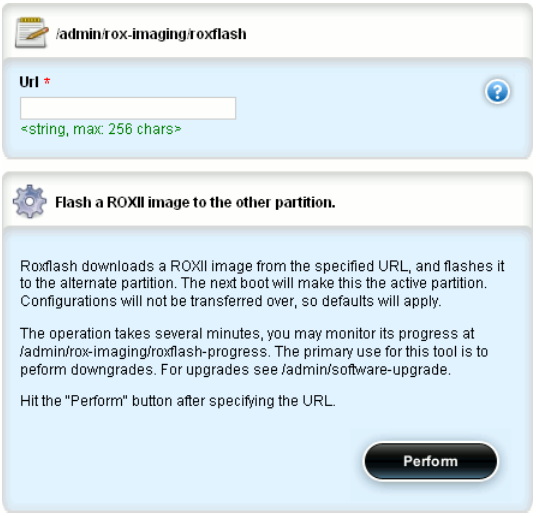
/admin/rox-imaging/roxflash
Url *
<string, max: 256 chars>
Flash a ROXII image to the other partition.
Roxflash downloads a ROXII image from the specified URL, and flashes it to the alternate partition. The next boot will make this the active partition. Configurations will not be transferred over, so defaults will apply.
The operation takes several minutes, you may monitor its progress at /admin/rox-imaging/roxflash-progress. The primary use for this tool is to peform downgrades. For upgrades see /admin/software-upgrade.
Hit the "Perform" button after specifying the URL.
Perform

Figure 2.33. ROXFlash forms

To perform a ROXFlash operation, enter the URL for the flash image into the ROXflash form and then click *Perform*. Next, monitor the progress by returning to the ROXflash Monitoring form located at *admin/rox-imaging*.

2.7. Scheduling Jobs

Use job scheduling to execute CLI (command line interface) commands at a specified time and date or in response to configuration changes. The path to the *scheduler* menu is *admin/scheduler*.

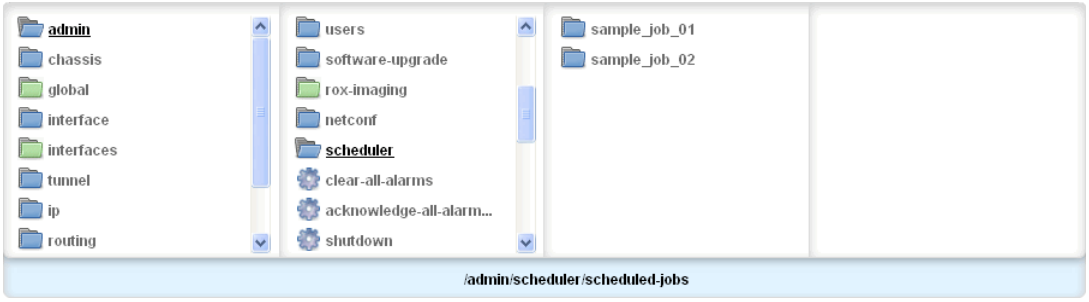

Figure 2.34. Scheduler menu

There are two types of scheduled jobs:

- *periodic* jobs launch at a defined interval. Set the interval in the *Minute*, *Hour*, *Day of Month*, and *Month* parameters. Use the *Day of Week* parameter to launch the job on a specific day of the week, such as every Friday. For information on how periodic scheduled jobs behave when you omit date and time parameters, see Figure 2.36, “Scheduled Jobs Form” and the field descriptions.
- *configchange* jobs launch only when the configuration changes.

The job scheduler *Command* parameter accepts most ROX CLI commands. Do not use commands that require a manual response or confirmation.

The */admin/scheduler/scheduled-jobs* table lists the scheduled jobs and their settings:

/admin/scheduler/scheduled-jobs
Job Name Job Type Minute Hour Day of Month Month Day of Week Command
sample_job_01 periodic 1 not found not found not found monday backupconfig
sample_job_02 periodic 5 5:00 1 1 monday clearmessagelog

Figure 2.35. Scheduled-jobs table

To add a scheduled job:

- Enter edit mode, navigate to *admin/scheduler*, and click **<Add scheduled-jobs>**.
- On the *Key settings* form, enter a name for the job and click **Add**.
- On the *Scheduled Jobs* form, set the job parameters.

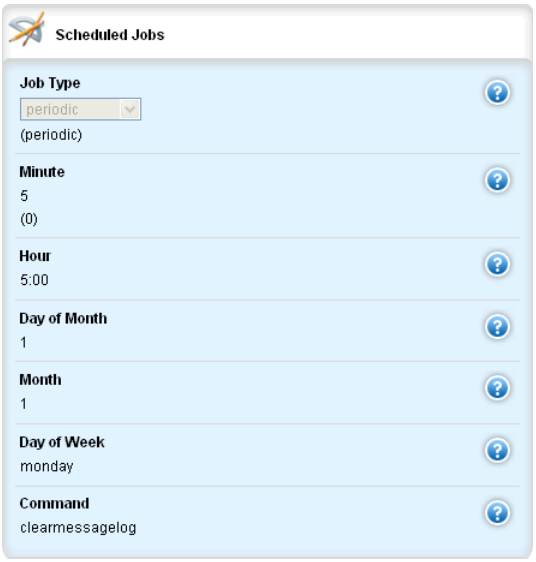

Figure 2.36. Scheduled Jobs form

Job Type

Synopsis: *string - one of the following keywords { periodic, configchange }*

Default: *periodic*

Determines when to launch the scheduled job:

- periodic: the job launches at a set date and time.
- configchange: the job launches when the configuration changes.

Minute

Synopsis: *A string*

Default:

For periodic jobs, sets the minutes portion of the job launch time. Valid values are in the range of 0 to 59. If no value is set, the scheduler uses the default value of 0 and launches the job every hour on the the hour.

- To specify a single value, enter the value in the field. For example, to launch the job 10 minutes past the hour, enter 10
- To specify a list of values, enter the values as a comma-separated list. For example, to launch the job at 14, 30, and 45 minutes past the hour, enter 15,30,45

- To specify a range of values, enter the range as comma-separated values. For example, to launch the job every minute between 30 and 45 minutes past the hour, enter 30-45

Hour

Synopsis: *A string*

For periodic jobs, sets the hour portion of the job launch time, in the 24-hour clock format. Valid values are in the range of 0 to 23. If no value is set, the job launches every hour at the time set in the Minute field.

- To specify a single value, enter the value in the field. For example, to launch the job at 5:00 pm, enter 17
- To specify a list of values, enter the values as a comma-separated list. For example, to launch the job at 9:00 am, 12:00 pm, and 5:00 pm, enter 9,12,17
- To specify a range of values, enter the range as comma-separated values. For example, to launch the job every hour between 9:00 am and 5:00 pm, enter 9-17

Day of Month

Synopsis: *A string*

For periodic jobs, sets the day of the month on which to run the scheduled job. Valid values are in the range of 1 to 31. If no value is set, the job launches every day.

- To specify a single value, enter the value in the field. For example, to launch the job on the tenth day of the month, enter 10
- To specify a list of values, enter the values as a comma-separated list. For example, to launch the job on the first, fifteenth, and thirtieth days of the month, enter 10,15,30
- To specify a range of values, enter the range as comma-separated values. For example, to launch the job on days one through fifteen, enter 1-15

Month

Synopsis: *A string*

For periodic jobs, sets the month in which to run the scheduled job. Valid values are in the rage of 1 to 12. If no value is set, the job launches every day.

- To specify a single value, enter the value in the field. For example, to set the month to February, enter 2
- To specify a list of values, enter the values as a comma-separated list. For example, to set the months to January, June, and December, enter 1,6,12
- To specify a range of values, enter the range as comma-separated values. For example, to set the months to January through June, enter 1-6

Day of Week

Synopsis: *A string*

For periodic jobs, sets the day of the week on which to run the scheduled job. Valid entries are in the range of 0 to 6, where 0 represents Sunday, 1 represents Monday, and so on. If no value is set, the job launches every day.

- To specify a single value, enter the value in the field. For example, to set the day to Monday, enter 1
- To specify a list of values, enter the values as a comma-separated list. For example, to set the days to Friday, Saturday, and Sunday, enter 5,6,0
- To specify a range of values, enter the range as comma-separated values. For example, to set the days to Monday through Friday, enter enter 1-5

Command

Synopsis: *A string*

The CLI commands to execute at the scheduled time. The command or list of commands can be up to 1024 characters in length. For example, this command saves the running configuration to a file named 'myconfig': show running-config | save myconfig

Do not use interactive commands or commands that require a manual response or confirmation.

2.8. The Featurekey

2.8.1. Overview

Some ROX™ software features are only available by purchasing an appropriate feature level. Consult the product datasheet for available feature levels and the specific capabilities they enable.

When specifying a feature level at the time of ordering, the featurekey is entered into the electronic signature on the device . The featurekey is independent of the compact flash card and is retained by the device should the card be replaced.

2.8.2. Upgrading Feature Levels in the field

Feature levels can be purchased and upgraded in the field with a file-based featurekey. To update your featurekey, contact your RuggedCom sales representative. For RX15xx products, you need to provide the serial number for the unit you are upgrading. The upgraded featurekey is licensed for the serial number you provide. For instructions on how to view your serial numbers, see Section 2.8.4, “Viewing RuggedCom Serial Numbers”.

To install the featurekey file, use the Install Files form found under that *admin* menu. You can also use the **file scp-featurekey-from-url** command from the ROX™ Command Line Interface. For instructions on how to upload the featurekey file, see Section 2.8.5, “Uploading a Featurekey”.

The upgraded featurekey resides on the device’s compact flash card. ROX™ evaluates both the device featurekey and the file-based featurekey, and then enables the most capable feature level described by the keys.

When using file-based featurekeys, the feature level follows the compact flash card. Moving the compact flash card to another device moves the feature level to the new device. If you want the upgraded feature level to be tied to a specific device, contact your RuggedCom sales representative to arrange for an RMA (Return to Manufacturer Authorization) to have the featurekey programmed into the device.

2.8.3. When a File-based featurekey does not Match the Hardware

In rare circumstances, you may need to remove the compact flash card from one device and transfer it to another device. For example: you may have a backup device to replace a malfunctioning unit, and you choose to use the upgrade featurekey on the malfunctioning unit’s compact flash card to retain your configuration in the backup unit.

The file-based featurekey on the compact flash card is licensed for a particular unit, but can be transferred to another unit to ensure continuity of service. When you transfer the file-based featurekey from its licensed unit to another unit for which it is not licensed, the device behaves in the following manner:

1. The device enables the higher feature level found on the compact flash card.
2. The device raises a non-clearable alarm, indicating a hardware mismatch with the featurekey.
3. The alarm trips the fail-safe relay and turns on the main alarm LED.

To acknowledge the alarm and resolve the issue, follow these steps:

1. Acknowledge the alarm. (For instructions on acknowledging alarms , see Chapter 6, *Alarms*.)
2. Contact a Ruggedcom sales representative and order a featurekey matching the serial numbers of the hardware you are using.

2.8.4. Viewing RuggedCom Serial Numbers

When you order a new featurekey, you need to provide RuggedCom with the chassis serial number. This section describes how to view your device's serial numbers through the **CLI** screen in the ROX™ web interface.

Follow these steps to display the serial numbers for your device:

Procedure 2.1. Viewing RuggedCom Serial Numbers

1. Launch a web browser and navigate to your device's IP address. Log in to ROX™. The ROX web interface appears.
2. Click the **Tools** tab and click the **CLI** link. The CLI screen appears.

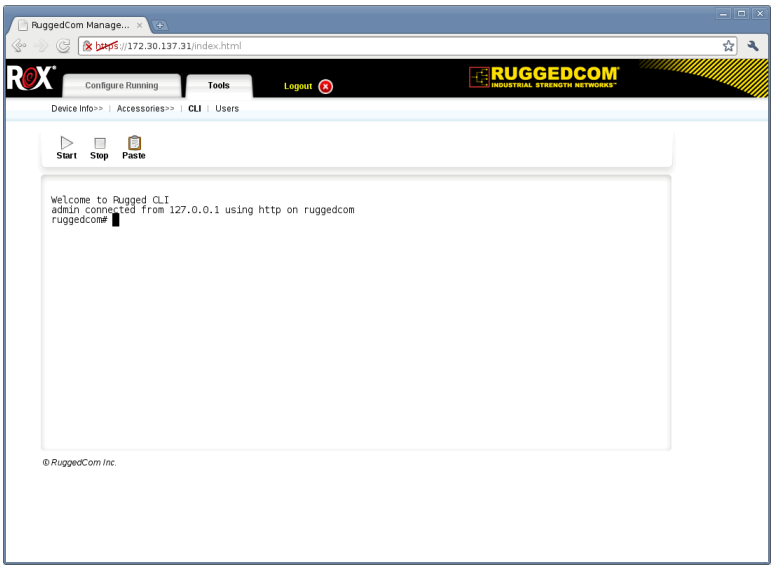

Figure 2.37. CLI in the ROX™ Web Interface

3. At the Operational mode command line prompt, type **show chassis** and press **Enter**. Chassis information appears:

```
ruggedcom# show chassis
chassis
 chassis-status
  model RX1501 software license "Layer 2 Standard Edition" order code ...
 hardware
  slot-hardware
      ORDER
SLOT  FIELD  DETECTED MODULE                                          SERIAL NUMBER
-------------------------------------------------------------------------------------
pm1   XX     none                                                     none
lm1   XX     none                                                     none
lm2   TC4    T1/E1 w/ 4x RJ48                                         L15R-3333-PR301
lm3   D02    DDS w/ 1x RJ48                                           7
lm4   XX     none                                                     none
lm5   CG01   1000TX w/ 2x RJ45                                        L15R-3109-PR001
lm6   XX     none                                                     none
main  CM04A  RX1501 8 Gigabit Layer 2 w/ 6 LM slots and 1 PM slots    R15R-1310-PR032
```

In the **`slot-hardware`** table, make note of the **`main`** slot serial number (highlighted in **bold text** in the example above).

4. When ordering a new featurekey, provide the **`main`** slot serial number to RuggedCom.

2.8.5. Uploading a Featurekey

After receiving your featurekey file from RuggedCom, save the file to a computer that is accessible to your device through your network.

2.8.5.1. Uploading a Featurekey Using the Web User Interface

Install Featurekey files using the Install Files forms found under the *admin* menu.

To install a featurekey file, navigate to *admin/install-files*. The Install Files form appears. In the in the *File type* field, select **featurekey**. In the *URL* field, enter the URL to the file. On the *Install Files to Devices* form, click the *Perform* button.

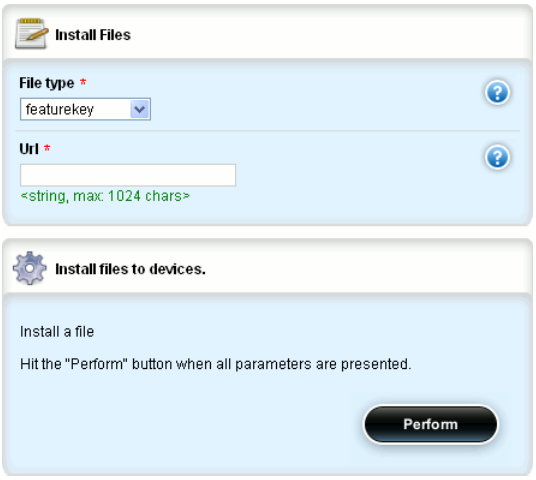

Figure 2.38. Install Files forms

For more information on installing files, see Section 2.9.1, “Installing Files”.

2.8.5.2. Uploading a Featurekey Using the Command Line Interface

To upload the file to your device, you will need to know the following information:

- the featurekey filename.
- a user name and password to log in to the computer where you saved the featurekey file.
- the hostname or IP address of the computer where you saved the featurekey file.

Follow these steps to upload a featurekey file to your device:

Procedure 2.2. Uploading a Featurekey File

1. Launch a web browser and navigate to your device’s IP address. Log in to ROX™. The ROX™ web interface appears.
2. Click the **Tools** tab and click the **CLI** link. The CLI screen appears.
3. In Operational mode, at the command line prompt, type the following command:

   file scp-featurekey-from-url {username}@{host}:{path}{filename} {filename}

   Where:

   - **{username}** is the name of a user who can log in to the computer where you saved the featurekey file.
   - **{host}** is hostname or IP address of the computer where you saved the featurekey file.
   - **{path}** is the directory path to the featurekey file on the computer.

- **{filename}** is the name of the featurekey file.

For example:

file scp-featurekey-from-url wsmith@10.200.20.39:/files/keys/1_cmRX1K-12-11-0015.key 1_cmRX1K-12-11-0015.key

4. Type the command with your parameters and press **Enter**. When prompted, type the user's password and press **Enter**. The system uploads the featurekey file:

```
ruggedcom# file scp-featurekey-from-url wsmith@10.200.20.39:/files/keys/
1_cmRX1K-12-11-0015.key 1_cmRX1K-12-11-0015.key
wsmith@10.200.20.39's password:
1_cmRX1K-12-11-0015.key               100%  192     0.2KB/s   00:00
ruggedcom#
```

5. To view the contents of the featurekey file, type the following command:

file show-featurekey {filename}

Where:

- **{filename}** is the name of the featurekey file.

For example:

file show-featurekey 1_cmRX1K-12-11-0015.key

6. Type the command with your featurekey filename and press **Enter**. The system displays the contents of the featurekey file:

```
ruggedcom# file show-featurekey 1_cmRX1K-12-11-0015.key
GPG_FEATUREKEY_LEVEL=1
GPG_FEATUREKEY_CM_SERIALNUMBER=RX1K-12-11-0015
GPG_FEATUREKEY_SIGNATURE=iEYEABECAAYFAk091pAACgkQP2pya+G5kdZeKACeKdHUB2G1T73Dymq8IjSdYDKAiskAn
3abBpCEhfLXxY2ZlVbvGNwDZow2
ruggedcom#
```

7. On the CLI screen, click **Stop** to close the CLI session.

2.8.6. Backing Up a Featurekey Using the Web User Interface

Featurekey files can be backed up using the following forms. These forms are accessible from the *admin* menu.

To back up a featurekey file, navigate to *admin/backup-files*. The Backup Files form appears. In the *File type* field, select **featurekey**. Enter additional parameters on the form. On the *Backup Files From Devices* form, click the *Perform* button.

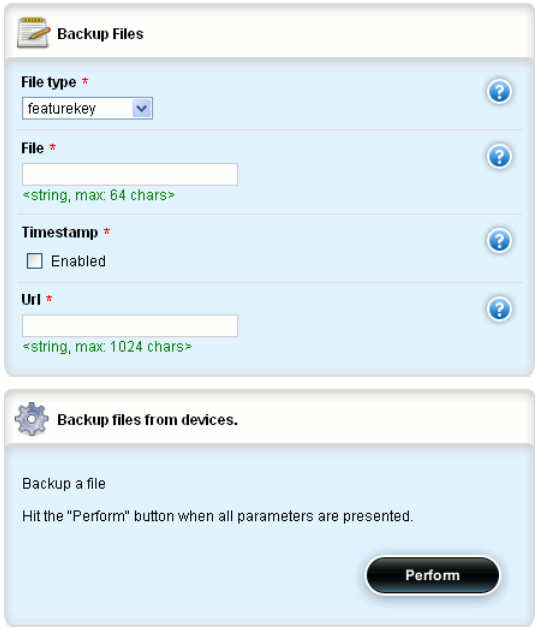

Figure 2.39. Backup Files forms

For more information on backing up files, see Section 2.9.2, “Backing Up Files”.

2.9. Installing and Backing Up Files

You can install and back up files using the following forms found under the *admin* menu.

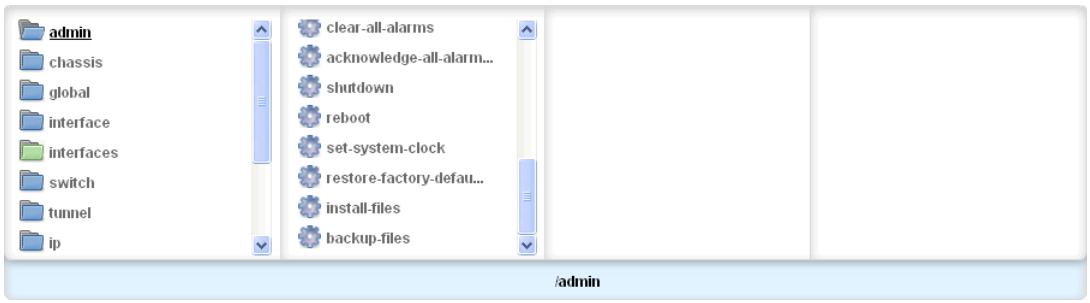

Figure 2.40. Administration menu

2.9.1. Installing Files

To install a file, click *install-files*. The Install Files forms appear.

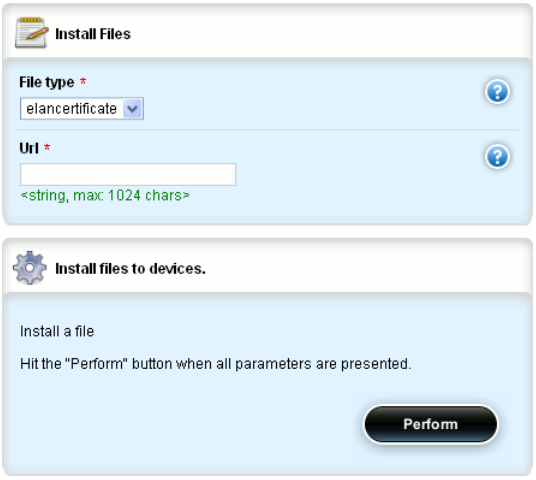

Figure 2.41. Install Files forms

On the *Install Files* form, select the file type and enter a URL. On the Install Files To Devices form, click the *Perform* button.

2.9.2. Backing Up Files

To back up a file, click on *backup-files*. The Backup Files forms appear.

Figure 2.42. Backup Files forms

On the *Backup Files* form, select the file type and enter the required parameters. On the Backup Files From Devices form, click the *Perform* button.

2.10. Deleting Log Files

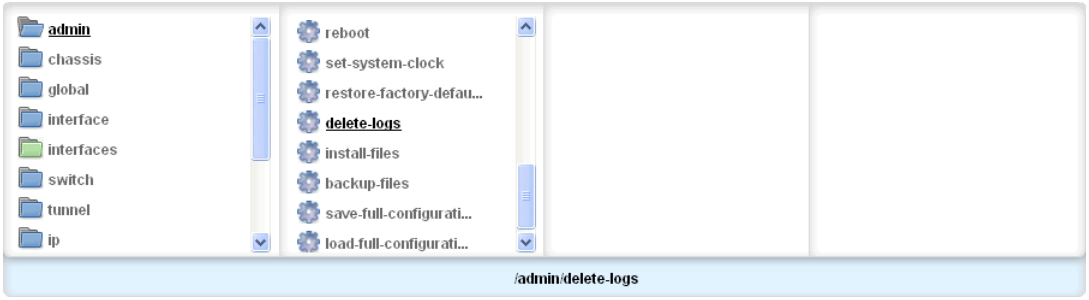

Figure 2.43. Delete-logs menu

To delete log files, click the **Perform** button on the Delete Log Files form. This form is accessible at *admin/delete-logs*.

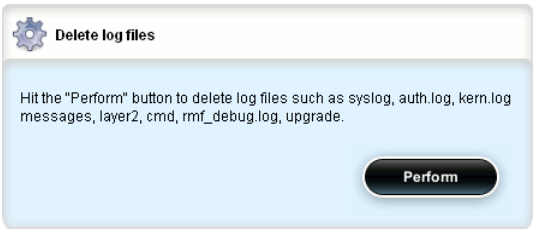
Delete log files
Hit the "Perform" button to delete log files such as syslog, auth.log, kern.log messages, layer2, cmd, rmf_debug.log, upgrade.
Perform

Figure 2.44. Delete Log Files form

2.11. Saving Full Configurations

Save full configurations to a file using the forms below. These forms are accessible at *admin/save-full-configuration*.

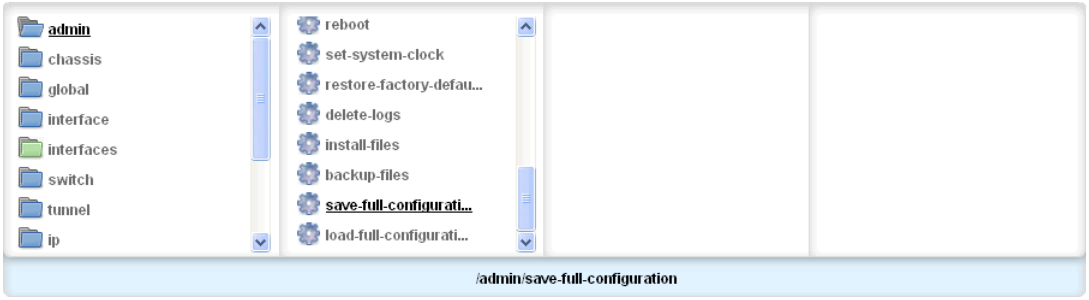

Figure 2.45. Save-full-configuration menu

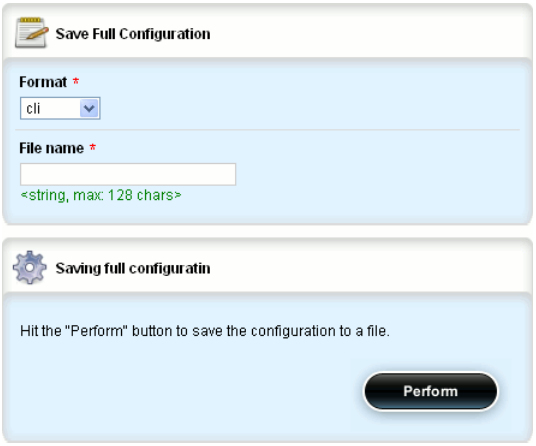

Figure 2.46. Save Full Configuration forms

To save full configurations to a file, select the format and enter the parameters in the Save Full Configuration form, then click the **Perform** button in the Saving Full Configuration form.

2.12. Loading Full Configurations

Load full configurations to a file using the forms below. These forms are accessible at *admin/load-full-configuration*.

Figure 2.47. Load-full-configuration menu

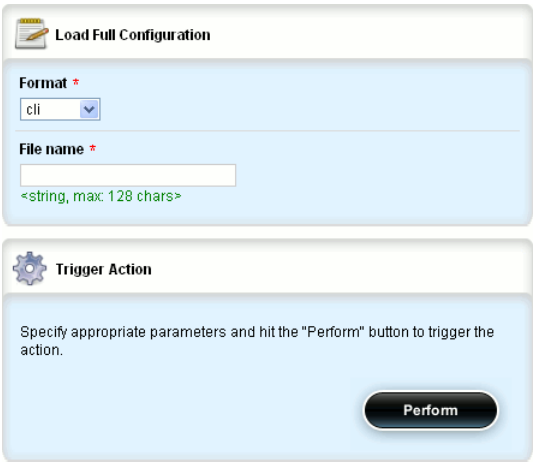

Figure 2.48. Load Full Configuration forms

To load full configurations to a file, select the format and enter the parameters in the Load Full Configuration form, then click the **Perform** button in the Trigger Action form.

3. Time Synchronization

ROX™ offers the following timekeeping and time synchronization features:

- Local hardware timekeeping and time zone management
- NTP time synchronization

3.1. NTP Fundamentals

NTP (Network Time Protocol) is an Internet protocol used to synchronize the clocks of computers to some time reference. Variants of NTP such as SNTP (Simple NTP, a reduced functionality NTP) and XNTP (Experimental NTP) exist. NTP itself is available in versions 3 and 4 (RuggedBackbone™ includes version 4).

NTP is a fault-tolerant protocol that allows an NTP daemon program to automatically select the best of several available time sources, or reference clocks, to synchronize to. Multiple candidates can be combined to minimize the accumulated error. Temporarily or permanently wrong time sources are detected and avoided.

The NTP daemon achieves synchronization by making small and frequent changes to the router hardware clock.

The NTP daemon operates in a *client-server mode*, both synchronizing from *servers* and providing synchronization to *peers*.

If NTP has a number of servers to choose from, it will synchronize with the lowest stratum server. The stratum is a measure of the number of servers to the (most highly accurate) reference clock. A reference clock itself appears at stratum 0. A server synchronized to a stratum n server will be running at stratum n + 1.

You will generally configure lower stratum NTP hosts as servers and other NTP hosts at the same stratum as peers. If all your configured servers fail, a configured peer will help in providing the NTP time. It is generally a good idea to configure one at least one server and peer.

The NTP daemon will know about the NTP servers and peers to use in three ways.

- It can be configured manually with a list of servers to poll,
- It can be configured manually with a list of peers to send to,
- It can look at advertisements issued by other servers on multicast or broadcast addresses.

Note that if multicasting or broadcasting is used, it is strongly recommended to enable authentication unless you trust all hosts on the network.

NTP uses UDP/IP packets for data transfer because of the fast connection setup and response times UDP offers. The NTP protocol uses port UDP port 123. Note that if your router employs a firewall and acts as a client it must open UDP port 123. Additionally, if the router acts as a server the firewall must allow connection requests on port 123 as well.

3.1.1. The NTP Sanity Limit

The NTP daemon corrects the system time through two means, “stepping” and “slewing”. If the difference between the local clock and the reference clock chosen by NTP (the “offset”) is more than 128ms for a period of more than 900 seconds, NTP will “step” or instantaneously correct the time. If the time difference is less than 128ms, NTP will “slew” the time by no more than 500 microseconds every second towards the correct time, in such a way that to an application on the system, the time never appears to be flowing backwards.

After booting, NTP uses slewing to achieve synchronization by making small and frequent changes to the router hardware clock. If the reference server's clock differs from the local clock by more than 1000 seconds, the NTP daemon decides that a major problem has occurred and terminates.

3.2. Configuring Time Synchronization

To configure time synchronization, configure the following items:

- set the system time and date. See Section 3.2.1, "Configuring the System Time and Date".
- set the system timezone. See Section 3.2.2, "Configuring the System Time Zone".
- set the local time settings. See Section 3.2.3, "Configuring the Local Time Settings".
- add remote NTP servers. You can add remote NTP servers with or without authentication. See Section 3.2.4, "Configuring NTP Servers".
- set the NTP server restrictions. See Section 3.2.6, "Configuring NTP Server Restrictions".
- configure an NTP server using Multicast or Broadcast. See Section 3.2.7, "Configuring an NTP Server using Multicast or Broadcast".
- configure an NTP client using Multicast. See Section 3.2.8, "Configuring an NTP Client using Multicast".
- configure an NTP client using Broadcast. See Section 3.2.9, "Configuring an NTP Client using Broadcast".

After configuring NTP, you can check the status of the NTP service. See Section 3.2.10, "Checking NTP Status".

3.2.1. Configuring the System Time and Date

To set the system time and date:

- Navigate to *admin/set-system-clock*.
- On the *Set New Time and Date* form, enter the date in the format YYYY-MM-DD HH:MM:SS.

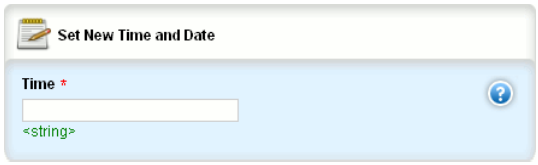

Figure 3.1. Set new Time and Date form

- On the *Set clock on target device* form, click **Perform**.

3.2.2. Configuring the System Time Zone

To set the system time zone:

- In edit mode, navigate to *admin*.
- On the *Timezone* form, select a timezone from the list.

  The Etc/GMT timezones conform to the POSIX style and have their signs reversed from common usage. In POSIX style, zones west of GMT have a positive sign; zones east of GMT have a negative sign.

Figure 3.2. Timezone form

- Commit the changes.

3.2.3. Configuring the Local Time Settings

On the *Local Time Settings* form, you enable the local clock and set the NTP stratum level.

The path to the *Local Time Settings* form is */services/time/ntp*.

To set the local time settings:

- In edit mode, navigate to */services/time/ntp*.
- On the *Local Time Settings* form, set the local time parameters.
- Commit the changes.

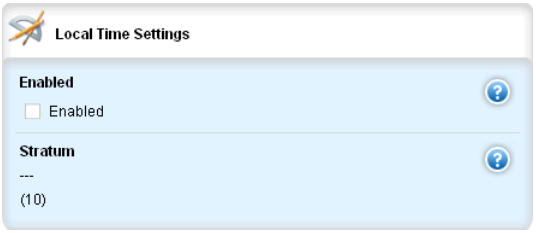

Figure 3.3. Local Time Settings form

Enable

Enables the local clock

Stratum

Synopsis: *unsigned byte integer*

Default: *10*

The stratum number of the local clock

3.2.4. Configuring NTP Servers

ROX™ can periodically refer to an NTP server to correct any accumulated drift in the onboard clock. ROX™ can also serve time via SNTP to hosts that request it.

You can add NTP servers with or without authentication keys. To associate an authentication key with an NTP server, you must first define the server key. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To view the list of configured NTP servers, navigate to */services/time/ntp/server*.

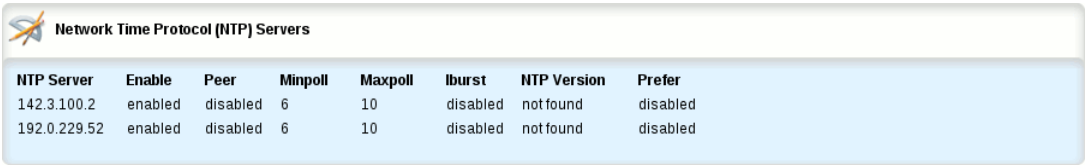
Network Time Protocol (NTP) Servers
NTP Server Enable Peer Minpoll Maxpoll Iburst NTP Version Prefer
142.3.100.2 enabled disabled 6 10 disabled not found disabled
192.0.229.52 enabled disabled 6 10 disabled not found disabled

Figure 3.4. Network Time Protocol (NTP) Servers

To add an NTP server:

- In edit mode, navigate to */services/time/ntp/server* and click **<Add server>**.
- On the *Key settings* form, enter the IP address or hostname for the server and click **Add**.
- On the *Network Time Protocol (NTP) Servers* form, set the server parameters.
- Commit the changes.

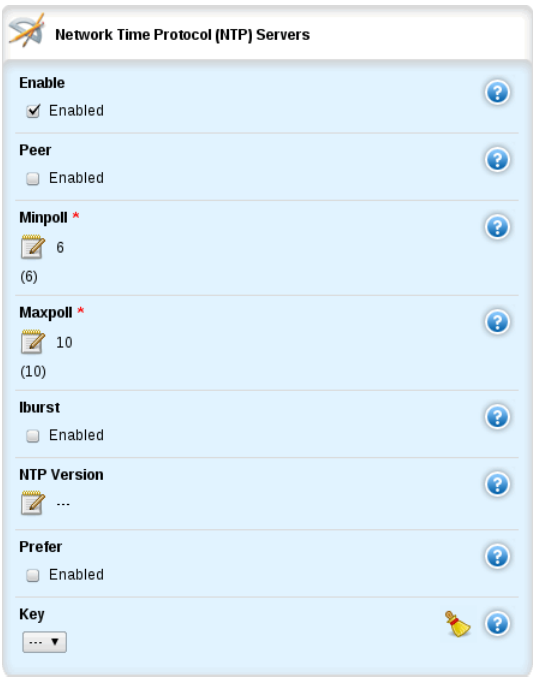

Figure 3.5. Network Time Protocol (NTP) Servers form

Enable

Turns on the NTP interface to this server.

Peer

Allows you to enter and edit peers. Peers are NTP servers of the same stratum as the router, and are useful when contact is lost with the hosts in the NTP servers menu.

Minpoll

Synopsis: *unsigned byte integer*

Default: *6*

Minimum poll interval for NTP messages, in seconds as a power of two.

Maxpoll

Synopsis: *unsigned byte integer*

Default: *10*

Maximum poll interval for NTP messages, in seconds as a power of two.

Iburst

When the server is unreachable and at each poll interval, send a burst of eight packets instead of the usual one.

NTP Version

Synopsis: *integer*

The version of the NTP protocol used to communicate with this host. Change this only if it is known that the host requires a version other than 4.

Prefer

Marks this server as preferred.

Key

Synopsis: *unsigned short integer*

An authentication key associated with this host.

3.2.5. Adding Server Keys

Use server keys to use authentication for NTP communications. NTP authentication authenticates the time source to help prevent tampering with NTP timestamps. When using authentication, both the local and remote servers must share the same key and key identifier. Packets sent to and received from the server/peer include authentication fields encrypted using the key.

Keys defined here are associated with NTP servers on the *Network Time Protocol (NTP) Servers* and *NTP Broadcast/Multicast Servers* forms.

To add a server key:

- In edit mode, navigate to */services/time/ntp/key* and click **<Add key>**.
- On the *Key settings* form, enter an identifier for the key and click **Add**.
- On the *Server Keys* form, set the key parameters.
- Commit the changes.

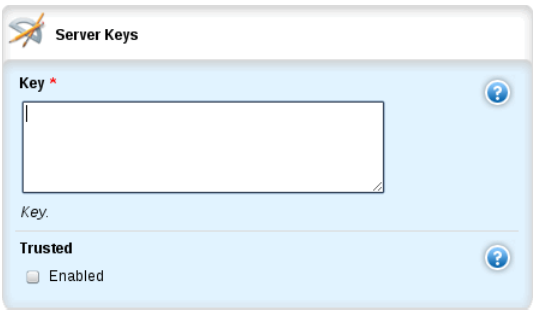

Figure 3.6. Server Keys form

Key

Synopsis: *"AES CFB128"-encrypted string*

Key.

Trusted

Mark this key is trusted for the purposes of authenticating peers with symmetric key cryptography. The authentication procedures require that both the local and remote servers share the same key and key identifier.

3.2.6. Configuring NTP Server Restrictions

Use server restrictions to control and restrict access to the NTP server.

To set NTP server restrictions:

- In edit mode, navigate to */services/time/ntp/restrict* and click **<Add restrict>**.
- On the *Key settings* form, set the following parameters and click **Add**.

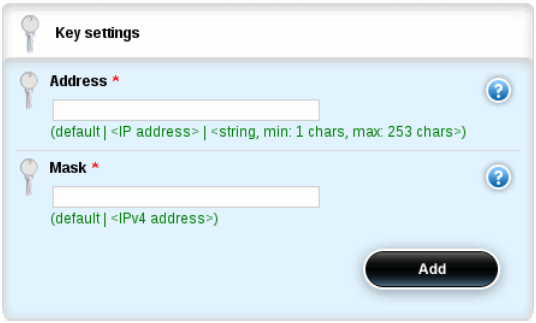

Figure 3.7. Server Restrictions Key settings form

Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Synopsis: *Domain name (RFC 1034)*

Synopsis: *string - the keyword { default }*

Address to match. The address can be host or network IP address or a valid host DNS name.

Mask

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *string - the keyword { default }*

Mask used to address match. Mask 255.255.255.255 means address is treated as the address of an individual host.

- On the *Server Restrictions* form, set the restriction parameters.
- Commit the changes.

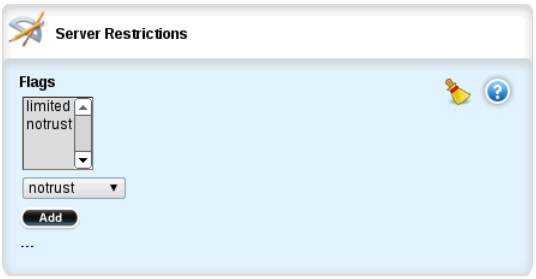

Figure 3.8. Server Restrictions form

Flags

Synopsis: *string - one of the following keywords { version, ntpport, notrust, notrap, noserve, noquery, nopeer, nomodify, lowpriotrap, limited, kod, ignore }*

Synopsis: *"flags" occurs in an array.*

Flags restrict access to NTP services. An entry with no flags allows free access to the NTP server.

- version: denies packets that do not match the current NTP version.
- ntpport: matches only if the source port in the packet is the standard NTP UDP port (123).
- notrust: denies service unless the packet is cryptographically authenticated.
- notrap: declines to to provide mode 6 control message trap service to matching hosts.
- noserve: denies all packets except ntpq(8) and ntpdc(8) queries.
- noquery: denies ntpq(8) and ntpdc(8) queries.

- nopeer: denies packets which result in mobilizing a new association.
- nomodify: denies ntpq(8) and ntpdc(8) queries attempting to modify the state of the server; queries returning information are permitted.
- lowpriotrap: declares traps set by matching hosts to be low priority.
- limited: denies service if the packet spacing violates the lower limits specified in the NTP discard setting.
- kod: sends a kiss-o-death (KoD) packet when an access violation occurs.
- ignore: denies all packets.

3.2.7. Configuring an NTP Server using Multicast or Broadcast

The NTP broadcast/multicast address must be the same as the client address. It is recommended that NTP authentication be used and that a server key be set with the broadcast/multicast setting. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To set a multicast/broadcast address for an NTP server:

- In edit mode, navigate to */services/time/ntp/broadcast* and click **<Add broadcast>**.
- On the *Key settings* form, enter the broadcast/multicast IP address and click **Add**.
- On the *NTP Broadcast/Multicast Servers* form, set the broadcast/multicast parameters.
- Commit the changes.

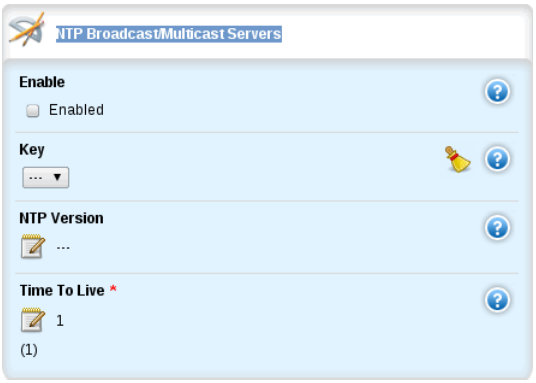

Figure 3.9. NTP Broadcast/Multicast Servers form

Enable

Enables sending broadcast or multicast NTP messages to this address.

Key

Synopsis: *unsigned short integer*

Authentication key.

NTP Version

Synopsis: *integer*

The version of the NTP protocol used to communicate with this host. Change this only if it is known that the host requires a version other than 4.

Time To Live

Synopsis: *unsigned byte integer*

Default: *1*

Time to live.

3.2.8. Configuring an NTP Client using Multicast

Configuring a multicast address for an NTP client enables the client to listen for and receive NTP messages on the multicast address. It is recommended that NTP authentication be used and that a server key be set with the multicast setting. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To set a multicast address for an NTP client:

- In edit mode, navigate to */services/time/ntp*.
- On the *NTP Multicast Clients* form, set the multicast parameters.
- Commit the changes.

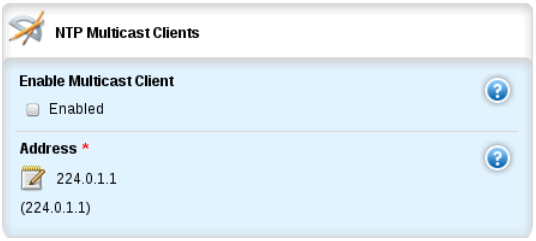

Figure 3.10. NTP Multicast Clients form

Enable Multicast Client

Enables the multicast message mode

Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Synopsis: *Domain name (RFC 1034)*

Default: *224.0.1.1*

The multicast address on which the NTP client listens for NTP messages.

3.2.9. Configuring an NTP Client using Broadcast

Configuring a broadcast address for an NTP client enables the client to listen for and receive NTP messages on the broadcast address, and enables the NTP server to send NTP messages on the broadcast/multicast address. It is recommended that NTP authentication be used and that a server key be set with the broadcast setting. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To set a broadcast address for an NTP client:

- In edit mode, navigate to */services/time/ntp*.
- On the *Network Time Protocol (NTP)* form, set the broadcast parameters.
- Commit the changes.

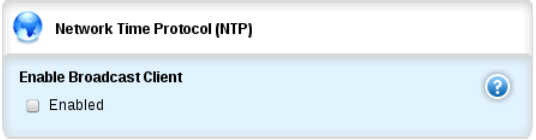

Figure 3.11. Network Time Protocol (NTP) form

Enable Broadcast Client

The broadcast address on which the NTP client listens for NTP messages.

3.2.10. Checking NTP Status

To view the NTP service status:

- In normal or edit mode, navigate to */services/time/ntp/ntp-status* and click <ntp-status>.
- On the *Trigger Action* form, click **Perform**.
- Review the NTP service status in the *NTP Service Status* form.

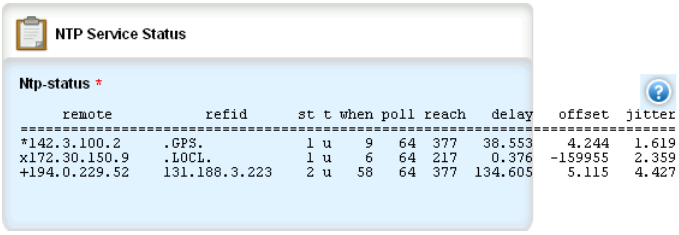

Figure 3.12. NTP Service Status form

For more information on viewing NTP status information, refer to http://support.ntp.org/bin/view/Support/TroubleshootingNTP

4. Basic Network Configuration

This chapter discusses the following:

- IP Interfaces
- Configuring IPv4 and IPv6 Addresses
- Simple Network Setups with IPv4 and IPv6 Addresses

4.1. IP Interfaces

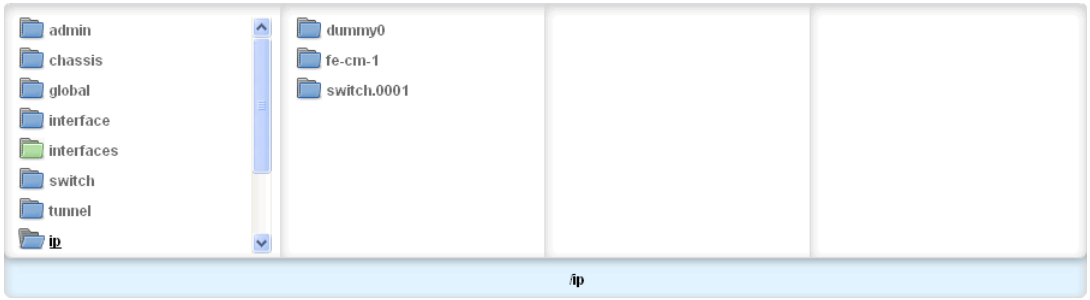

Figure 4.1. IP menu

The IP menu is accessible from the main menu under *ip*.

4.1.1. Configuring an IP Address

The RX1500 has the following internet interfaces configured by default: *dummy0*, *fe-cm-1*, and *switch.0001*. The default IP addresses for *fe-cm-1* and *switch.0001* are configured under the *ipv4* submenu. *switch.0001* is the VLAN interface and is only seen if you have one or more ethernet line modules. It is created implicitly as all switched ports have a default PVID of 1. The following table lists the default IP addresses.

Interface	IP Address
switch.0001	192.168.0.2/24
fe-cm-1	192.168.1.2/24

Table 4.1. Default IP Addresses

To configure a different IP address on an interface, see Procedure 4.1, “Configuring an IP Address”.

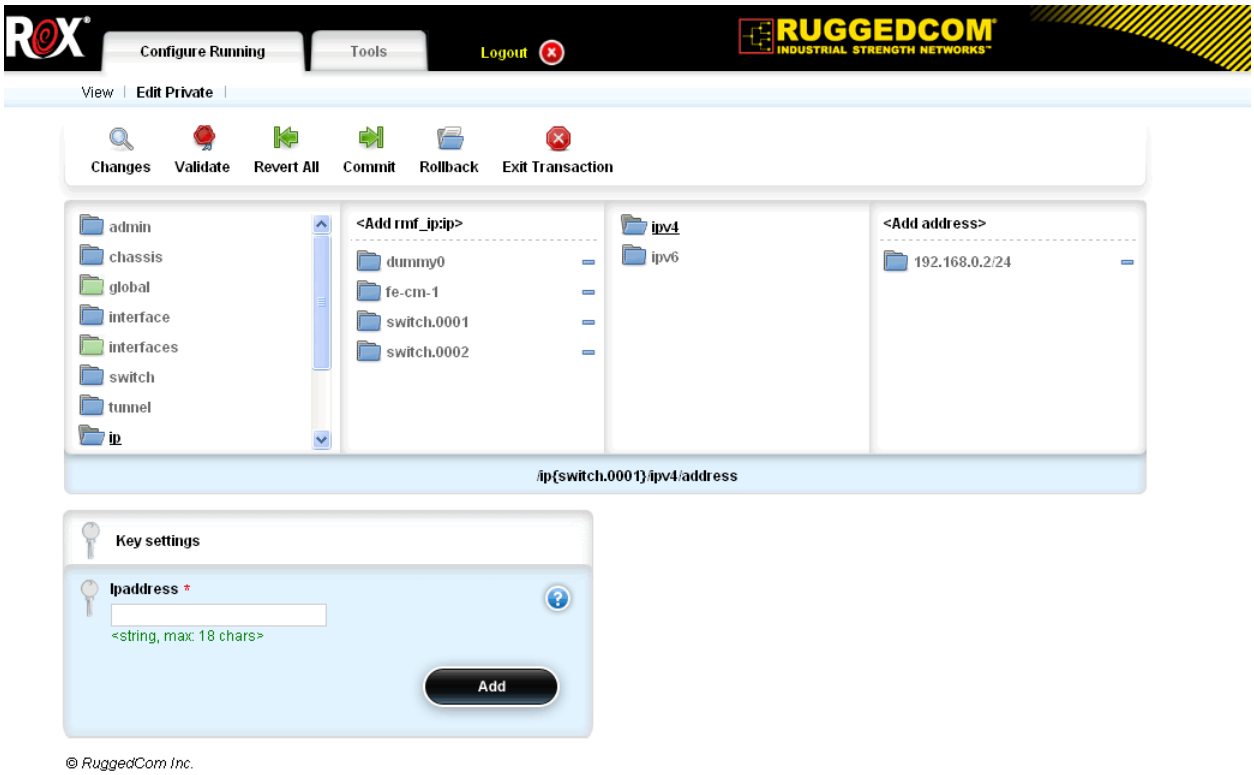

Figure 4.2. Configuring an IP Address

Procedure 4.1. Configuring an IP Address

1. Enter Edit Private mode.
2. Navigate to *ip/interface/ipv4*.
3. To delete an existing IP address, click the *delete* icon.
4. Click *Add address*. The Key settings form appears.
5. In the *IPaddress* field, type the new IP address.
6. Click **Commit**.
7. Click **Exit Transaction**.

To create additional interfaces, see Section 5.3, “Adding Interfaces to Switched Ports”.

4.1.2. Simple Network Setup with the Default IPv4 Addresses

This section describes how to set up a simple network using the factory default IPv4 address.

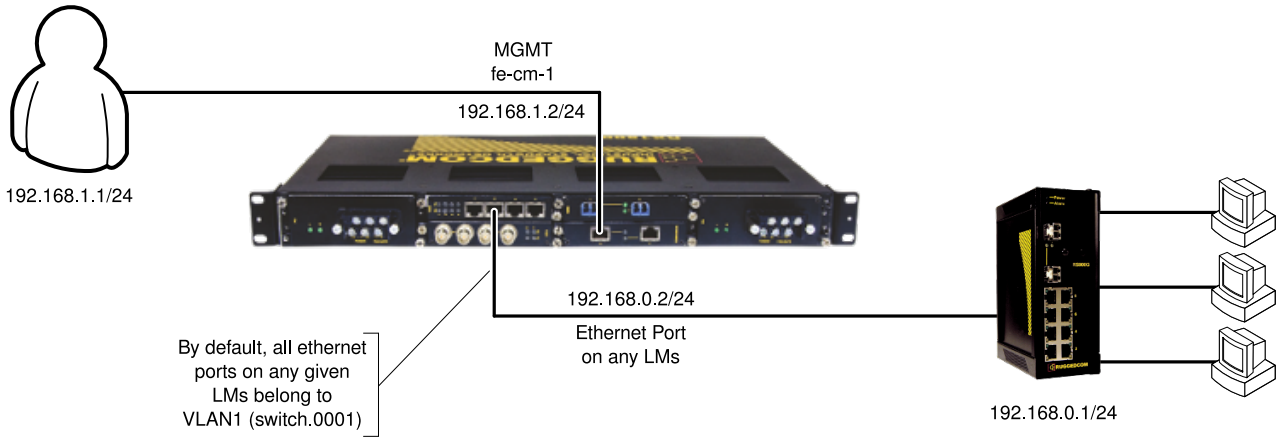

Figure 4.3. Basic Network Setup Using the Default IPv4 Addresses

Procedure 4.2. Basic Network Setup Using the Default IPv4 Addresses

1. Connect a user PC to the Fast Ethernet port (fe-cm-1) of the RX1500 and configure the PC to be on the same subnet as the port.
2. Configure the PC to use the IP address of the Fast Ethernet port as the default gateway
3. Connect one of the switched ports from any available LMs to a switch typically connecting a LAN
4. The PCs connected to the switch should be on the same subnet as the switch.
5. Configure the switch and the PCs behind the switch to use Switch.0001's IP address (192.168.0.2) as the default gateway
6. From the user PC, ping the IP addresses of the PCs behind the switch. Verify the ping is successful.

To configure a WAN port and assign an IP address, see Chapter 23, *WAN*.

To configure Dynamic Routing on the unit, see Chapter 34, *Dynamic Routing*.

To configure Static Routes and Default Gateways, see Chapter 35, *Static Routing*.

For information related to the Firewall and IP NAT that might be necessary before connecting the unit to the INTERNET, see Chapter 38, *Firewall*.

For information on adding VLAN interfaces to Switched Ports (Ethernet Ports on LMs and SM) and assigning IP addresses to configured VLANs to make them routable, see Section 5.3, "Adding Interfaces to Switched Ports".

For information on Dynamic IP address assignment and ProxyARP on switched and non-switched ports, see Section 5.3.1.1, "Configuring IP Address Source and ProxyARP for VLAN Interfaces" and Section 5.4.1, "Configuring IP Address Source and ProxyARP for Non-switched Interfaces".

4.1.3. Configuring an IPv6 Address

IPv6 link local addresses starting with the prefix FE80 are assigned to all routable Ethernet interfaces in the RX1500. The Link Local addresses are hidden in the Web UI but they are visible from the CLI (Command Line Interface) using the **show interfaces ip** command.

To advertise IPv6 link layer addresses to their neighbors on the same link, IPv6 Router Advertisement in IPv6 Neighbor Discovery must be enabled. For more information on IPv6 fundamentals and Neighbor Discovery, see Section 5.1, "IPv6 Fundamentals" and Section 5.2, "IPv6 Neighbor Discovery".

Procedure 4.3. Configuring an IPv6 Address

1. Enter Edit Private mode.

2. From the WEB UI Navigate to ip/interface/ipv6.
3. Click Add address. The Key settings form appears.
4. In the IPaddress field, type an IPv6 address with a network prefix
5. Click Commit.
6. Click Exit Transaction.
7. To delete an existing IPv6 address, click the delete icon under ip/interface/ipv6.
8. Refer to steps 3 to 7 to configure a new IPv6 address

4.1.4. Simple Network Setup with IPv6 Addresses

This section describes how to configure a simple network using the factory default IPv6 address.

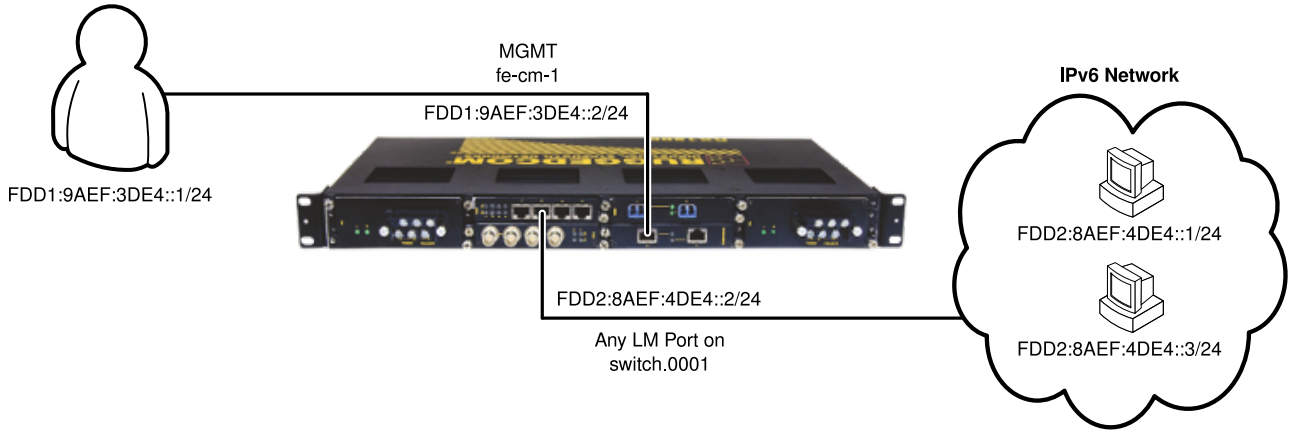

Figure 4.4. Simple IPv6 Network Setup

Procedure 4.4. Simple IPv6 Network Setup

1. Connect a user PC to Fast Ethernet port (fe-cm-1) of the RX1500 and configure the PC to be on the same subnet as the port.
2. Configure the S.PC with IPv6 address FDD1:9AEF:3DE4::1/24 and Default Gateway as FDD1:9AEF:3DE4::2.
3. Configure the fe-cm-1 and switch.0001 interfaces of the RX1500 with the IPv6 addresses shown in Figure 4.4, “Simple IPv6 Network Setup”.
4. Connect one of the switched ports from any available LMs to an IPv6 capable network.
5. Configure the D.PCs on the IPv6 network to be on the same IP subnet as switch.0001 and configure the Default Gateway address as FDD2:8AEF:4DE4::2/48.
6. Enable IPv6 Neighbor Discovery under ip/{interface}/ipv6/nd. For more information on IPv6 neighbor discovery, see Section 5.2, “IPv6 Neighbor Discovery”.
7. Confirm that you can reach the D.PCs from the S.PC.

4.1.5. Routable Interfaces

Figure 4.5. Routable Interfaces table

The Routable Interfaces table is accessible from the *ip* menu.

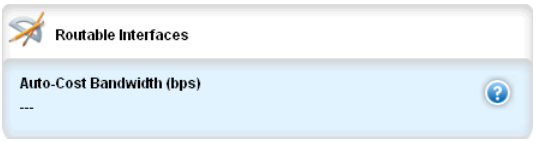

Figure 4.6. Routable Interfaces form

The path to the Routable Interfaces form is *ip/{interface}*.

Interface Name

Synopsis: *A string*

The name for this routable logical interface

Auto-Cost Bandwidth (kbps)

Synopsis: *unsigned long integer*

This value is used in auto-cost calculations for this routable logical interface in kbps

Figure 4.7. Addresses table

The path to the Addresses table is *ip/{interface}/ipv4*. The Addresses table provides a summary of which IP addresses are configured.

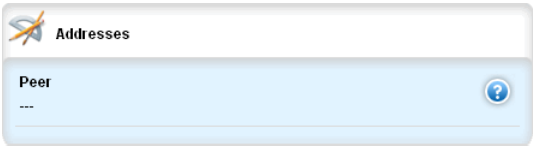

Figure 4.8. Addresses form

The path to the Addresses form is *ip/{interface}/ipv4/{address}*.

ipaddress

Synopsis: *IPv4 address and prefix in CIDR notation*

The IPv4/Prefix (xxx.xxx.xxx.xxx/xx).

peer

Synopsis: *IPv4 address in dotted-decimal notation*

The peer IPv4 Address (xxx.xxx.xxx.xxx, PPP link only).

5. IP Network Interfaces

This chapter familiarizes the user with:

- IPv6 Fundamentals and IPv6 Neighbor Discovery
- Adding VLAN Interfaces to Switched Ports
- Configuring IP Address Source and ProxyARP for Switched and Non-switched Interfaces

5.1. IPv6 Fundamentals

Version 6 of the Internet Protocol (IPv6, RFC 2460) has been designated to replace IPv4 throughout the Internet. Some important changes that IPv6 introduces relative to IPv4 fall into the following categories:

5.1.1. Addressing

IPv6 addresses are four times the length of IPv4 addresses, at 128 bits, to be used as 64 bits of network and 64 bits of host address. The larger address space allows much greater flexibility in hierarchical network definition and routing.

The IPv6 packet header has been simplified relative to IPv4 in order to simplify and therefore speed the processing of packets by routing nodes. It also features more efficiently encoded options and greater flexibility in creating extensions.

5.1.2. Security

Security has been designed into IPv6, rather than being treated as a component that must be added to existing IPv4 network stacks.

5.1.3. IPv6 Address Scopes

There are three scopes of IPv6 addresses named Link Local, Unique Local and Global. A Link Local address is automatically assigned to any IPv6 capable interface. This address is mandatory for the devices on the same link to communicate with each other.

The link local address begins with “FE80” in the first 10 bits of an IPv6 address and the address is not routable. The scope for Unique Local address is within enterprise networks. It identifies the boundary of private networks within an organization. Example of a link local address: **FE80:0000:0000:0000:020A:DCFF:FE01:0CCD**

Unique Local addresses are similar to private IPv4 addresses and they are not routable on the Internet. A Unique Local address consists of the first 7 bits as the site address starts with “FD”, the next 1 bit set to 1 meaning locally assigned, next 40 bits as the Global ID to identify a company, next 16 bits as the Subnet ID to identify the subnets within a site and it is usually defined based on hierarchical plan, and finally the last 64 bits for the Interface ID. Example of a unique local address: **FD00:ABAB:CDCD:EFEF:020A:DCFF:FE01:0CCD**

The Global IPv6 addresses are routable and they are interned to be used on the Internet. In order to allow address aggregation the global addresses are structured in hierarchical order. A global address is identified by the first 48 bits specified by the service provider as the global routing prefix in which the first 3 bits of the address start with 001 (2000::/3), the next 16 bits after the global routing prefix are used to define subnets and the last 64 bits are used for Interface ID to define a host. Example of a unique local address: **2001:0CCD:3456:789A:8A9C:BCAB:023A:1234**

5.1.4. IPv6 Multicast Addresses

In IPv6 multicast addresses are widely used. The use of broadcast address is removed in IPv6, instead IPV6 multicast addresses are used for neighbor discovery and route advertisement. An IPv6 multicast address starts with first 8 bits all set to 1 (FF), next 4 bits to define the Lifetime (0 - Permanent, 1 -

Temporary), then the following 4 bits to define the scope (1 - Node, 2 - Link, 5 - Site, 8 – Organization and E – Global) and the last 112 bits identify a multicast Group ID. Some well-known multicast addresses are mentioned below:

IPv6 M.Cast Address	Scope	Description
FF02::1	Link-Local	All Nodes on a Link
FF02::2	Link-Local	All Routers on a Link
FF01::1	Node-Local	Same Node
FF01::2	Node-Local	Same Router
FF05::2	Site-Local	All Routers on a Site
FF02::1:FFxx:xxxx	Link-Local	Solicited Node Address

Table 5.1. Multicast Addresses

5.2. IPv6 Neighbor Discovery

In IPv6 the Neighbor Discovery (ND) protocol is seen as a replacement for IPv4 ARP message. It uses ICMPv6 messages with various purposes include finding a link-layer address of a neighbor, discover neighbor routers, determine any change in the link-layer address, determine when a neighbor is down, send network information from router to hosts, which includes hop limit, MTU size, determining the network prefix used on a link, address auto configuration, and the default route information.

There many types of ICMPv6 messages among which five types of messages are used by the ND protocol. The five types of ICMPv6 messages are briefly described in the following section:

- Router Solicitation (ICMPv6 type 133): This message is sent by hosts to routers as a request to router advertisement message. It uses a destination multicast address: FF02::2
- Router Advertisement Messages (ICMPv6 type 134): This message is used by routers to announce its presence in a network. The message includes network information related to IPv6 prefixes, default route, MTU size, hop limit and auto configuration flag. It uses a destination multicast address: FF02::1
- Neighbor Solicitation Messages (ICMPv6 type 135): This message is sent by hosts to determine the existence of another host on the same. The goal is to find the link-layer of neighbor nodes on the same link.
- Neighbor Advertisement Messages (ICMPv6 type 136): This message is sent by hosts to indicate the existence of the host and it provides information about its own link-layer address.
- Redirect Messages (ICMPv6 type 137): This message is sent by a router to inform a host about a better router to reach a particular destination address.

In RX1500, Neighbor Discovery should be configured on all Ethernet interfaces enabled for IPv6. The following figure displays the available configuration options for IPv6 Neighbor Discovery.

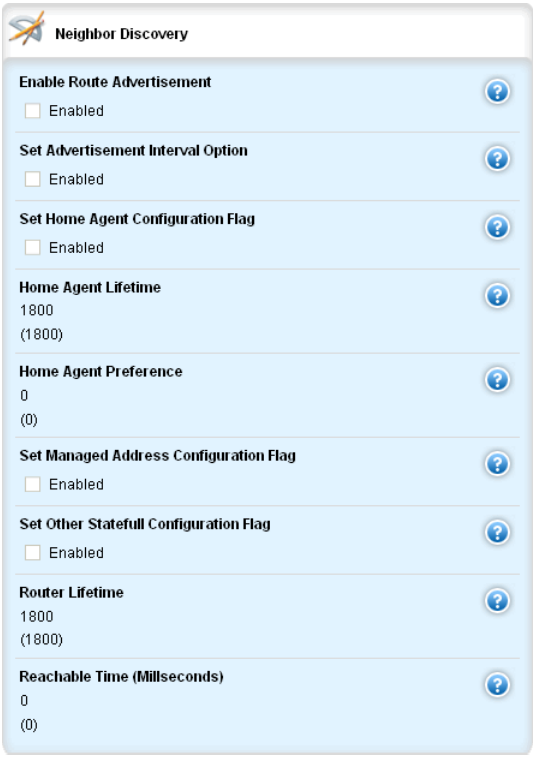

Figure 5.1. Neighbor Discovery form

The path to the Neighbor Discovery form is *ip/{interface}/ipv6/nd*.

Enable Route Advertisement

Enable to send router advertisement messages.

Set Advertisement Interval Option

Includes an Advertisement Interval option which indicates to hosts the maximum time in milliseconds, between successive unsolicited router advertisements.

Set Home Agent Configuration Flag

Sets/unsets the flag in IPv6 router advertisements which indicates to hosts that the router acts as a home agent and includes a home agent option.

Home Agent Lifetime

Synopsis: *unsigned integer*

Default: *1800*

The value to be placed in the home agent option, when the home agent config flag is set, which indicated the home agent lifetime to hosts. A value of 0 means to place a router lifetime value.

Home Agent Preference

Synopsis: *unsigned integer*

Default:

The value to be placed in the home agent option, when the home agent config flag is set, which indicates the home agent preference to hosts.

Set Managed Address Configuration Flag

The flag in IPv6 router advertisements, which indicates to hosts that they should use the managed (stateful) protocol for addresses autoconfiguraiton in addition to any addresses autoconfigured using stateless address autoconfiguration.

Set Other Statefull Configuration Flag

The flag in IPv6 router advertisements, which indicates to hosts that they should use the administered (stateful) protocol to obtain autoconfiguration information other than addresses.

Router Lifetime

Synopsis: *unsigned integer*

Default: *1800*

The value (in seconds) to be placed in the Router Lifetime field of router advertisements sent from the interface. Indicates the usefulness of the router as a default router on this interface. Setting the value to zero indicates that the router should not be considered a default router on this interface. It must be either zero or between the value specified with the IPv6 nd ra-interval (or default) and 9000 seconds. The default is 1800 seconds.

Reachable Time (Millseconds)

Synopsis: *unsigned integer*

Default:

The value (in milliseconds) to be placed in the Reachable Time field in the router advertisement messages sent by the router. The configured time enables the router to detect unavailable neightbors. The value zero means unspecified (by this router). The default is 0.

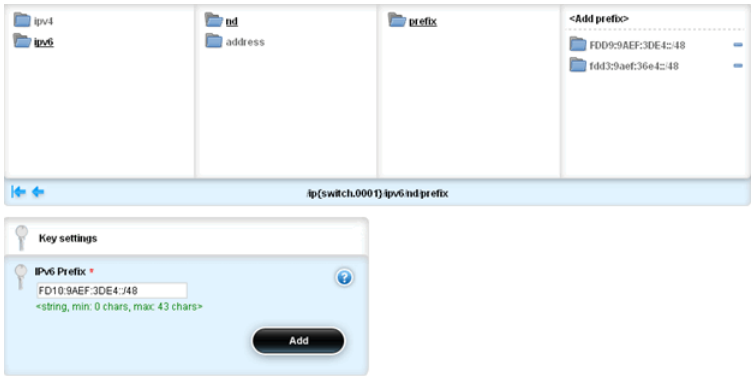

Figure 5.2. Neighbor Discovery IPv6 Prefix

An IPv6-capable interface can use Neighbor Discovery to advertise IPv6 network prefixes to its neighbor on the same link.

Figure 5.3. Neighbor Discovery IPv6 Prefix forms

IPv6 Prefix

Synopsis: *IPv6 address and prefix in CIDR notation*

The IPv6 network/prefix.

Valid Lifetime

Synopsis: *unsigned integer*

Synopsis: *string - the keyword { infinite }*

The length of time in seconds during what the prefix is valid for the purpose of on-link determination. The default value is 2592000.

Preferred Lifetime

Synopsis: *unsigned integer*

Synopsis: *string - the keyword { infinite }*

The length of time in seconds during which addresses generated from the prefix remain preferred. The default value is 604800.

Off Link

Indicates that advertisement makes no statement about on-link or off-link properties of the prefix.

No Autoconfig

Indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration.

Set Router Address Flag

Indicates to hosts on the local link that the specified prefix contains a complete IP address by setting the R flag.

This screen is accessible after adding an IPv6 Prefix under the Neighbor Discovery. To display the forms, navigate to *ip/{interface}/ipv6/nd/prefix*.

5.3. Adding Interfaces to Switched Ports

For switched ports, you create routable interfaces by configuring VLANs. VLANs are created either *implicitly* or *explicitly*. There are four locations in the web user interface where VLAN interfaces are created implicitly, and one location where they are created explicitly:

Explicit/Implicit	Location in the Web User Interface
Implicit	interface/switch/{port}
Implicit	interface/trunks
Implicit	switch/mcast-filtering/static-mcast-table
Implicit	switch/mac-table/static-mac-table
Explicit	switch/vlans/static-vlan

Table 5.2. Locations For Creating VLAN Interfaces

The procedure below is an example of how to create explicit VLAN interfaces.

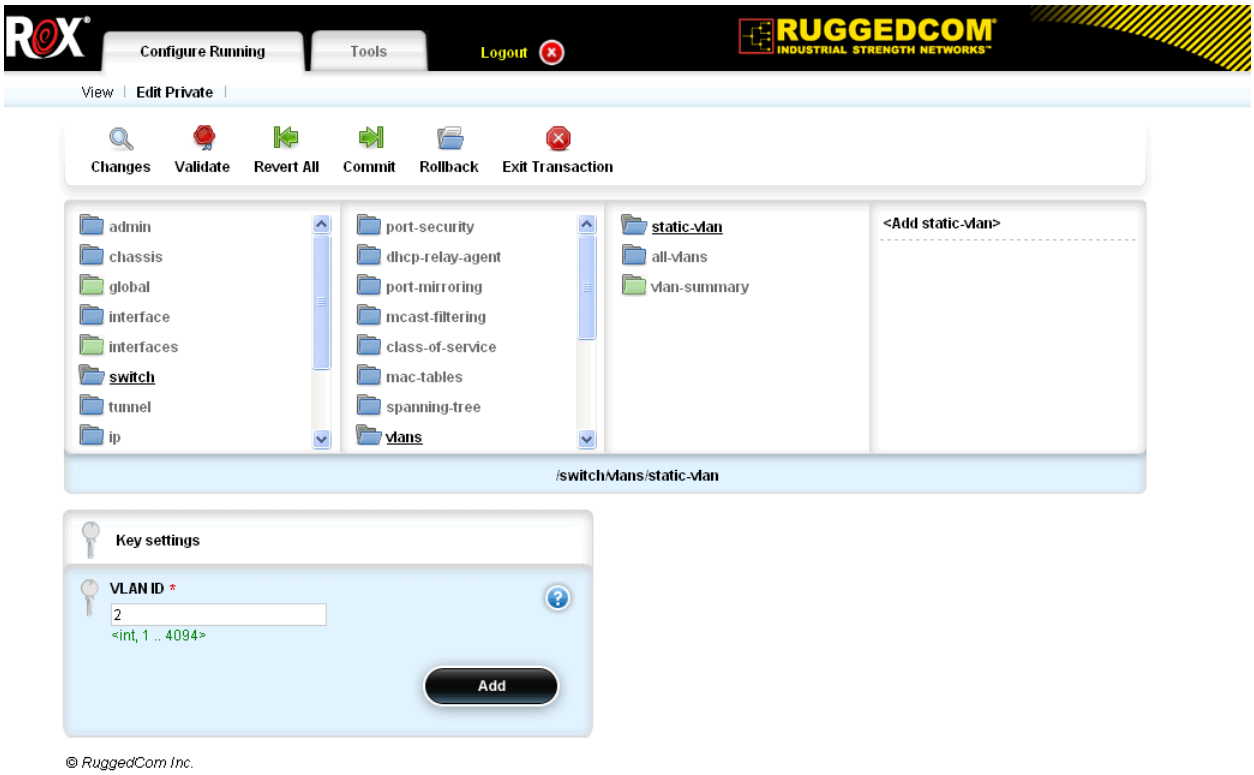

Figure 5.4. Explicitly Adding a VLAN Interface to a Switched Port

Procedure 5.1. Explicitly Adding a VLAN Interface at *switch/vlans/static-vlan*

1. Go into Edit Private mode.
2. Navigate to *switch/vlans/static-vlan*.
3. Click on *Add static-vlan*. The Key settings form appears.
4. In the *VLAN ID* field, enter a number from 1 to 4094 (for example, 2).
5. Click **Add**.
6. Click **Commit**.
7. Click **Exit Transaction**.

The procedures below are examples of how to create implicit VLAN interfaces.

Procedure 5.2. Implicitly Adding a VLAN Interface at *interface/switch/{port}*

1. Go into Edit Private mode.
2. Navigate to *interface/switch/{port}*. The switch forms are displayed.
3. On the *VLAN* form, type the PVID number into the PVID field.
4. Click **Commit**.
5. Click **Exit Transaction**.

Procedure 5.3. Implicitly Adding a VLAN Interface at *interface/trunks*

1. Go into Edit Private mode.
2. Navigate to *interface/trunks*.
3. Click on *Add trunks*. The Key settings form appears.

4. In the *Trunk ID* field, type a number between 1 and 15.
5. Click **Add**. The Trunks forms appear.
6. On the *VLAN* form, type a PVID number into the *PVID* field.
7. Click **Commit**.
8. Click **Exit Transaction**.

Procedure 5.4. Implicitly Adding a VLAN Interface at *switch/mac-tables/static-mac-table*

1. Go into Edit Private mode.
2. Navigate to *switch/mac-tables/static-mac-table*.
3. Click on Add static-mac. The Key settings form appears.
4. In the *MAC Address* field, type a string of 17 characters (for example, 11:22:33:44:55:66).
5. In the *VLAN ID* field, enter a number between 1 and 4094.
6. Click **Add**. The Static MAC Address Parameters form appears.
7. Click **Enabled** in the *Learned* field or select a port in the *Slot* field.
8. Click **Commit**.
9. Click **Exit Transaction**.

*When configuring the static-mac-table, you must click **Enabled** in the Learned field or select a port in the Slot field, otherwise the configuration will fail when you try to commit it.*

Procedure 5.5. Implicitly Adding a VLAN Interface at *switch/mcast-filtering/static-mcast-table*

1. Enter edit mode, navigate to *switch/mcast-filtering/static-mcast-table*, and click <Add static-mcast-table>. The *Key settings* form appears.
2. In the *VLAN ID* field, enter a number between 1 and 4094.
3. In the *MAC Address* field, type a string of 17 characters beginning with 01 (for example, 01:22:33:44:55:66).
4. Click **Add**. The *Static Multicast Summary* form appears. Select an option from the *CoS* field or leave **normal** as the default.
5. Click **Commit**.
6. Commit the changes.

ROX™ will create a new routable interface for each VLAN created (either implicitly or explicitly) on the switch. These interfaces have names such as "switch.xxxx" where "x" is the VLAN ID that has been created. It will not have a default IP address so you will need to create one using the procedure in Section 4.1, “IP Interfaces” or use DHCP. For more information on setting DHCP, see Section 5.4.1, “Configuring IP Address Source and ProxyARP for Non-switched Interfaces”.

5.3.1. All-VLANs

After VLAN interfaces have been added, they will be displayed in the All VLANs table, below. The path to this table is *switch/vlans/all-vlans*.

Figure 5.5. All VLANs table

5.3.1.1. Configuring IP Address Source and ProxyARP for VLAN Interfaces

The All VLANs Properties form can be used to configure ProxyARP and dynamic address source by following the procedures below.

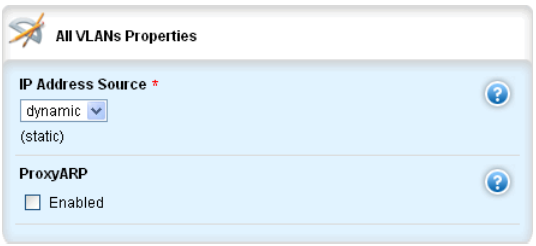

Figure 5.6. All VLANs Properties form

Procedure 5.6. Configuring IP Address Source and ProxyARP for VLAN Interfaces

1. Go into Edit Private mode.
2. Navigate to *switch/vlans/all-vlans/{vlan}*. The All VLANs Properties form is displayed.
3. In the *IP Address Source* field, select **dynamic** if you want the interface to get an IP address from a DHCP server. For information on configuring RX1500 as a DHCP server, see Chapter 15, *DHCP Server*. The default value for the IP Address Source field is **static**. To assign a static IP address to an interface, see Chapter 4, *Basic Network Configuration*.
4. Click **Commit**.
5. Click **Exit Transaction**.

Procedure 5.7. Configuring ProxyARP Using the All VLANs Properties form

1. Go into Edit Private mode.
2. Navigate to *switch/vlans/all-vlans/{vlan}*. The All VLANs Properties form is displayed.
3. In the *ProxyARP* field, click **Enabled**.
4. Click **Commit**.
5. Click **Exit Transaction**.

5.4. Non-switched Interface Menu

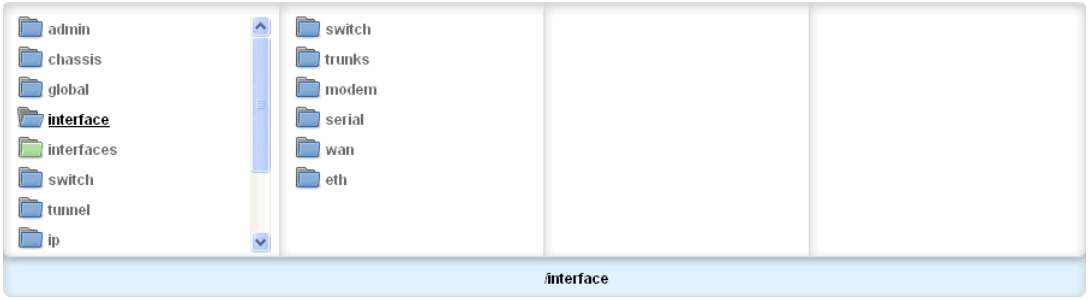

Figure 5.7. Non-switched Interface menu

The Non-switched (or Route-only) Interface menu is accessible from the main menu.

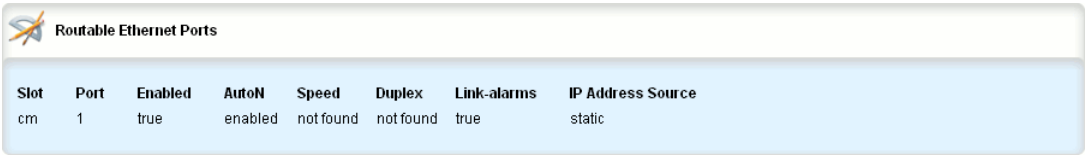
Routable Ethernet Ports
Slot Port Enabled AutoN Speed Duplex Link-alarms IP Address Source
cm 1 true enabled not found not found true static

Figure 5.8. Routable Ethernet Ports table

The path to the Routable Ethernet Ports table is *interface/eth*.

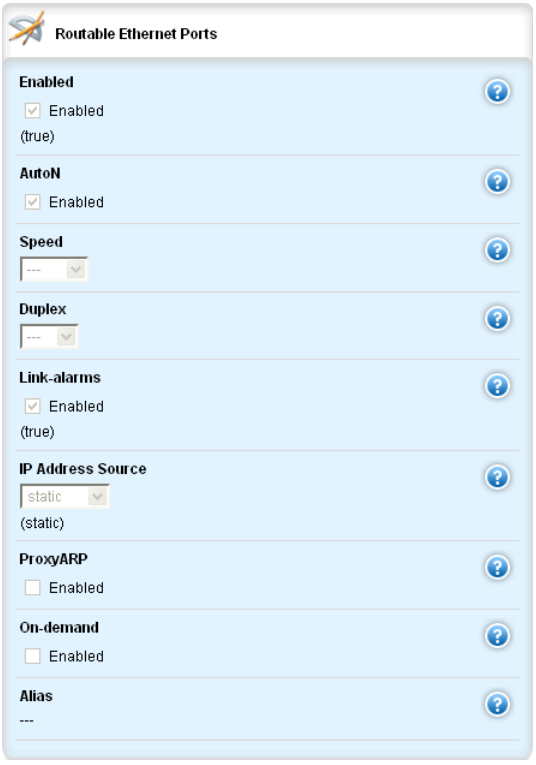

Figure 5.9. Routable Ethernet Ports form

The path to the Routable Ethernet Ports form is *interface/eth/{port}*.

Slot

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Enabled

Synopsis: *boolean*

Default: *true*

Enables/Disables the network communications on this port

AutoN

Enables or disables IEEE 802.3 auto-negotiation. Enabling auto-negotiation results in speed and duplex being negotiated upon link detection; both end devices must be auto-negotiation compliant for the best possible results

Speed

Synopsis: *string - one of the following keywords { 1000, 100, 10 }*

Speed (in Megabit-per-second or Gigabit-per-second). If auto-negotiation is enabled, this is the speed capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is explicitly forced to this speed mode. AUTO means advertise all supported speed modes.

Duplex

Synopsis: *string - one of the following keywords { full, half }*

If auto-negotiation is enabled, this is the duplex capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is explicitly forced to this duplex mode. AUTO means advertise all supported duplex modes.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP. Option DYNAMIC is a common case of dynamically assigned IP address. It switches between BOOTP and DHCP until it gets the response from the relevant server. Must be static for non-management interfaces

ProxyARP

Enables/Disables whether the port will respond to ARP requests for hosts other than itself

on-demand

This interface is up or down on demand of link fail over.

alias

Synopsis: *A string*

The SNMP alias name of the interface

5.4.1. Configuring IP Address Source and ProxyARP for Non-switched Interfaces

IP addresses on routable ports are static by default. To change the IP address of the port to dynamic, follow the procedure below. ProxyARP can also be enabled using this form.

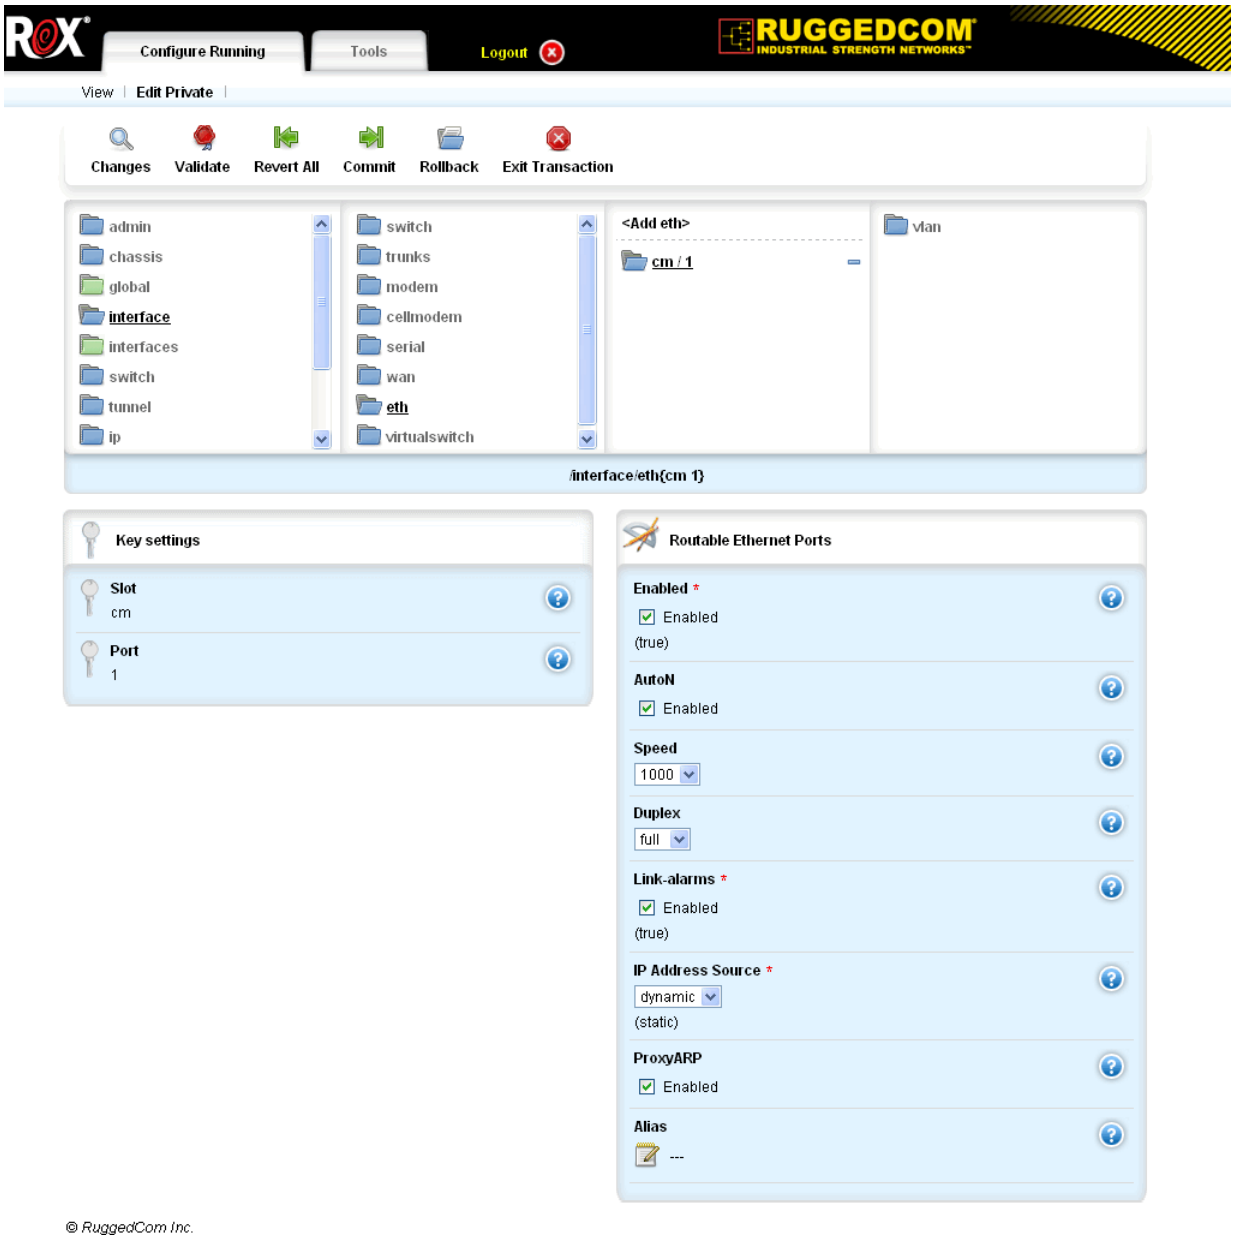

Figure 5.10. Configuring Dynamic Address Source and ProxyARP

Procedure 5.8. Configuring IP Address Source and ProxyARP for Non-switched Interfaces

1. Go into Edit Private mode.
2. Go to *interface/eth/{port}*. The Routable Ethernet Ports form appears.
3. In the *IP Address Source* field, select **dynamic** if you want the interface to get an IP address from a DHCP server. For information on configuring RX1500 as a DHCP server, see Chapter 15, *DHCP Server*. To assign a static IP address to an interface, see Chapter 4, *Basic Network Configuration*.

4. Click **Commit**.
5. Click **Exit Transaction**.

To set ProxyARP for a static or dynamic interface, follow the procedure below.

Procedure 5.9. Setting ProxyARP

1. Go into Edit Private mode.
2. Go to *interface/eth/(port}*. The Routable Ethernet Ports form appears.
3. In the *ProxyARP* field, click **Enabled**.
4. Click **Commit**.
5. Click **Exit Transaction**.

6. Alarms

6.1. Introduction

The ROXII alarm system is a highly configurable notification system of events of interest. Asserted alarms in the system may be viewed in a table in the CLI, web user interface, as well as queried by NETCONF. Alarms are categorized by subsystem.

The alarm system allows the user to:

- enable/disable alarms with the exception of mandatory alarms
- configure whether or not an alarm triggers the fail-relay and paints the alarm LED red
- configure the severity of an alarm to one of the following: emergency, alert, critical, error, warning, notice, info, debug (in descending order of severity). A small minority of alarms have fixed severity.

6.1.1. Alarm Subsystems

As of the current release, there are three subsystems that support alarms; they are Admin, Chassis, and Switch.

Note that some of the following examples describing the nature of each alarm subsystem may not be available in this release. A list of the available alarms can be viewed in the configuration file at **`/admin/alarm-cfg`**.

Admin Subsystem: these alarms are for administrative aspects of the device, including feature-key problems, upgrades, and configuration changes.

Chassis Subsystem: these alarms are for physical or electrical problems, or events of interest. This includes irregular voltages at the power supply or the insertion or removal of a module.

Switch Subsystem: these alarms pertain to layer-2 events of interests such as RSTP topology changes and link up/down events.

6.1.2. Fail-Relay Behavior

The fail-relay shall be activated when an active alarm in the system is also configured to trigger it. After an alarm has been acknowledged or cleared it ceases to assert the fail-relay. The fail-relay will only be de-activated when all active alarms that are configured to assert it have been acknowledged or cleared.

6.1.3. Alarm LED Behavior

The alarm LED on the control module shall be red when unacknowledged alarm(s) are asserted and the LED is enabled for any of the active alarms. After an alarm has been acknowledged or cleared, the LED is switched off.

6.1.4. Clearing and Acknowledging Alarms

There are two broad types of alarms:

1. Non-Clearable alarms - Users cannot clear these alarms, only acknowledge them; the difference between these actions is outlined later in this section. These alarms have a condition associated with them that the system assesses. The system asserts the alarm when the condition is true and clears the alarm when the condition has been resolved. An example of this is 'Bad input supply on power module'. If a redundant power module loses its supply an alarm is asserted. If the problem is resolved and power is returned to the module, the system de-asserts the alarm. De-asserted alarms remain as active alarms until acknowledged by the user.

2. Clearable alarms - these alarms simply report an event of interest that has no resolution per se. An example of this would be a 'configuration changed' alarm. These alarms are clearable by the user and are never cleared by the system.

Alarms may be cleared and acknowledged both on an individual basis and globally (i.e. clear/acknowledge all active-alarms). When an alarm is cleared by the user it is removed from the active alarms table and no longer asserts the fail-relay and LED. When an alarm is acknowledged by the user it de-asserts the fail-relay and LED, but it remains in the active alarms table, unless the alarm is non-clearable and de-asserted by the system. In the latter case it is removed from the table, because the condition was resolved.

6.2. Alarm Configuration

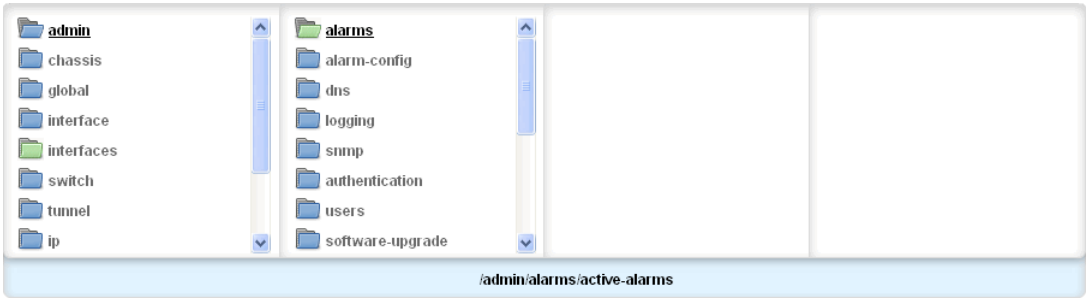

Figure 6.1. Alarms menu

The Alarms menu is accessible from the main menu under *admin*.

View active alarms in the Active Alarms table.

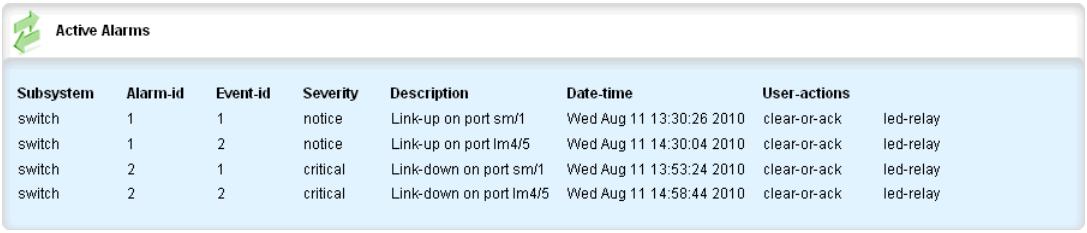
Active Alarms
Subsystem Alarm-id Event-id Severity Description Date-time User-actions
switch 1 1 notice Link-up on port sm/1 Wed Aug 11 13:30:26 2010 clear-or-ack led-relay
switch 1 2 notice Link-up on port lm4/5 Wed Aug 11 14:30:04 2010 clear-or-ack led-relay
switch 2 1 critical Link-down on port sm/1 Wed Aug 11 13:53:24 2010 clear-or-ack led-relay
switch 2 2 critical Link-down on port lm4/5 Wed Aug 11 14:58:44 2010 clear-or-ack led-relay

Figure 6.2. Active Alarms table

If data is configured, the Active Alarms table will appear on the same screen as the *Alarms* menu.

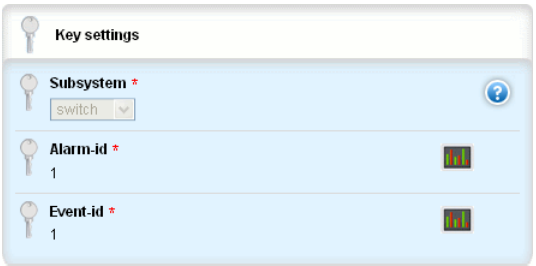

Figure 6.3. Active Alarms Key Settings form

If data is configured, the path to the Key Settings form and Active Alarms form is *admin/alarms/{interface}*.

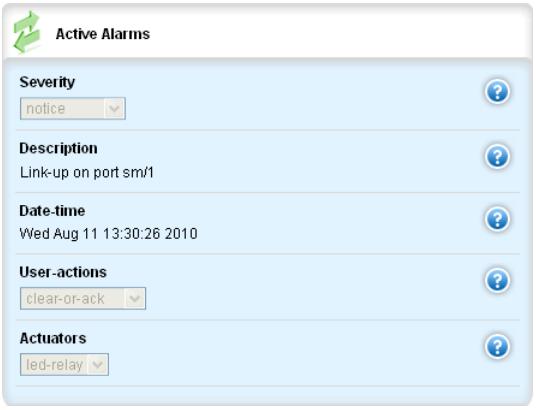

Figure 6.4. Active Alarms form

subsystem

Synopsis: *string - one of the following keywords { wan, switch, chassis, admin }*

Alarms are categorized by the subsystem to which they belong e.g.: Admin, Chassis, Ethernet, WAN.

Alarm ID

Synopsis: *integer*

Alarm Type Identifier. A value that uniquely defines a type of alarm.

Event ID

Synopsis: *integer*

Alarm Event Identifier. A value that uniquely defines a specific alarm event of the indicated alarm type.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The class of severity: Emergency, Alert, Critical, Error, Warning, Notice, Info, Debug

description

Synopsis: *string*

When applicable, provides further details on the alarmable event

Date/Time

Synopsis: *string*

The date and time the event was detected

User Actions

Synopsis: *string - one of the following keywords { must-resolve, clear-or-ack, resolve-or-ack }*

There are three categories of alarms:

1. clear or ack : the user can clear (remove from 'active-alarm' list) and/or acknowledge (turn off actuator(s) but keep as active-alarm).
2. ack or resolve : the user can acknowledge only, the system will clear the alarm once it is acknowledged and the condition is resovled.
3. must-resolve : for a minority of alarms, the condition must be resolved to turn off actuators and clear the alarm.

actuators

Synopsis: *string - one of the following keywords { acked, none, led-relay, led, relay }*

Indicates which actuator(s) this alarm currently asserts. 'ACKED' indicates the alarm was acknowledged so actuators are de-asserted.

Individual alarms can be cleared or acknowledged on the Clear Alarm Menu Action form or the Acknowledge Alarm Menu Action form. To clear or acknowledge an alarm, select *admin/alarms/{alarms submenu}* and then select the *Clear* action or the *Acknowledge* action.

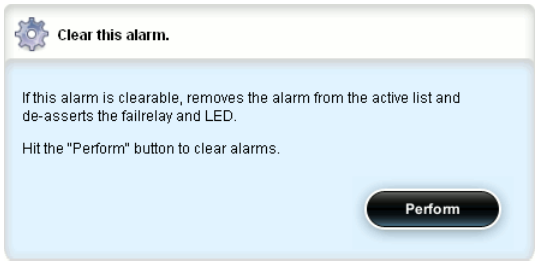
Clear this alarm.
If this alarm is clearable, removes the alarm from the active list and de-asserts the failrelay and LED.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.5. Clear Alarm Menu Action form

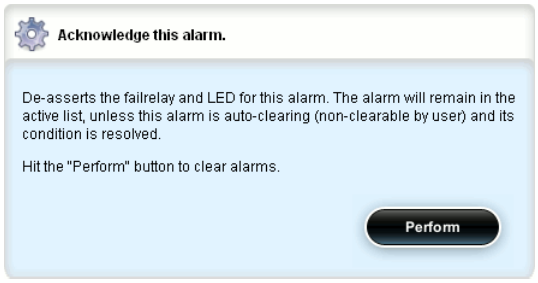
Acknowledge this alarm.
De-asserts the failrelay and LED for this alarm. The alarm will remain in the active list, unless this alarm is auto-clearing (non-clearable by user) and its condition is resolved.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.6. Acknowledge Alarm Menu Action form

To clear or acknowledge ALL alarms, instead of only individual alarms, access the Clear All Alarms and Acknowledge All Alarms menu action forms. These forms are accessible from the *admin* menu. The path to the Clear All Alarms Menu Action and the Acknowledge All Alarm Menu Action is *admin*, then clicking on the *clear-all-alarms* action or the *acknowledge-all-alarms* action.

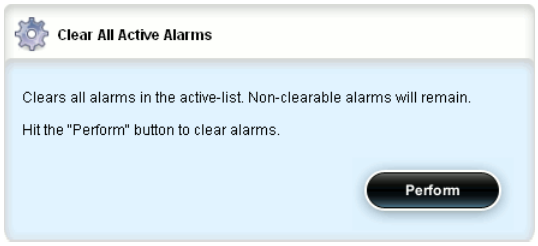
Clear All Active Alarms
Clears all alarms in the active-list. Non-clearable alarms will remain.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.7. Clear All Alarms Menu Action form

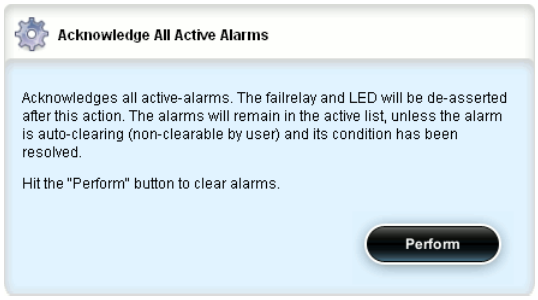
Acknowledge All Active Alarms
Acknowledges all active-alarms. The failrelay and LED will be de-asserted after this action. The alarms will remain in the active list, unless the alarm is auto-clearing (non-clearable by user) and its condition has been resolved.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.8. Acknowledge All Alarms Menu Action form

6.2.1. Administrative Alarm Configuration

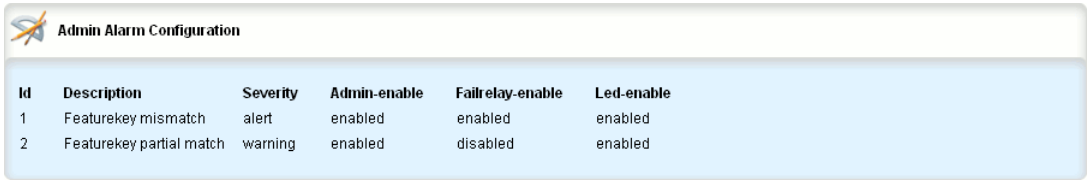
Admin Alarm Configuration
Id Description Severity Admin-enable Failrelay-enable Led-enable
1 Featurekey mismatch alert enabled enabled enabled
2 Featurekey partial match warning enabled disabled enabled

Figure 6.9. Admin Alarm Configuration table

The path to the Admin Alarm Configuration table is *admin/alarm-config/admin*.

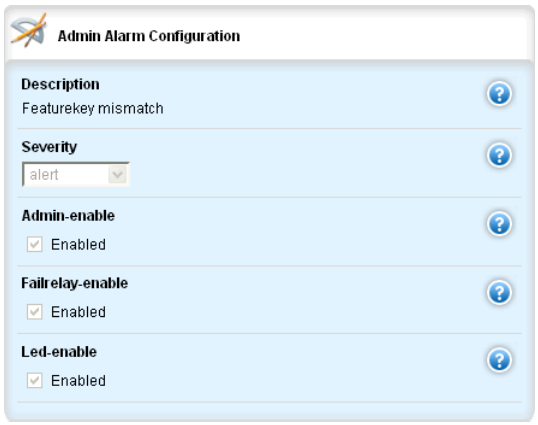

Figure 6.10. Admin Alarm Configuration form

The path to the Admin Alarm Configuration form is *admin/alarm-config/admin/{alarm id}*.

id

Synopsis: *integer*

This is the ID number of the alarm assigned by the system.

description

Synopsis: *A string*

The name of the alarm.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The severity level can be one of emergency, alert, critical, error, warning, notice, info, and debug. This cannot be changed for some alarms

admin-enable

If disabled, the alarm is not reported in the active list and does not actuate led/failrelay.

failrelay-enable

If enabled, this alarm will assert the failrelay.

led-enable

If enabled, the main 'Alarm' LED light will be red when this alarm is asserted. If disabled, the main 'Alarm' LED light is not affected by this alarm.

6.2.2. Chassis Alarm Configuration

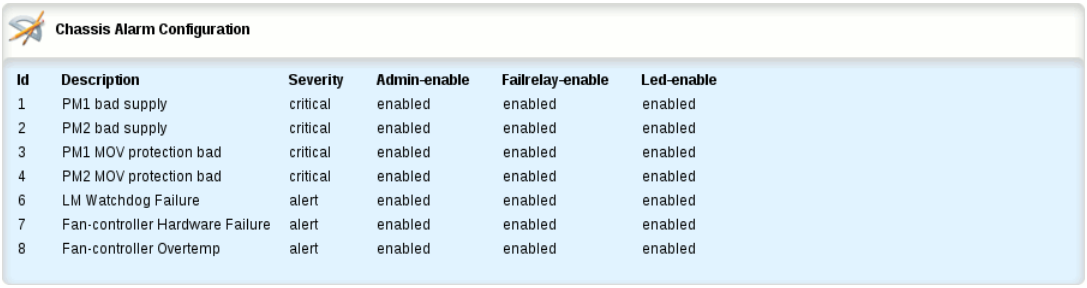

Chassis Alarm Configuration

Id	Description	Severity	Admin-enable	Failrelay-enable	Led-enable
1	PM1 bad supply	critical	enabled	enabled	enabled
2	PM2 bad supply	critical	enabled	enabled	enabled
3	PM1 MOV protection bad	critical	enabled	enabled	enabled
4	PM2 MOV protection bad	critical	enabled	enabled	enabled
6	LM Watchdog Failure	alert	enabled	enabled	enabled
7	Fan-controller Hardware Failure	alert	enabled	enabled	enabled
8	Fan-controller Overtemp	alert	enabled	enabled	enabled

Figure 6.11. Chassis Alarm Configuration table

The path to the Chassis Alarm Configuration form is *admin/alarm-config/chassis*.

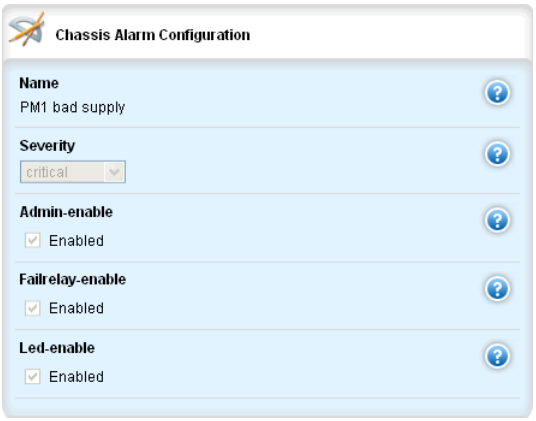

Figure 6.12. Chassis Alarm Configuration form

The path to the Chassis Alarm Configuration form is *admin/alarm-config/chassis/{alarm id)*.

id

Synopsis: *integer*

This is the ID number of the alarm assigned by the system.

description

Synopsis: *A string*

The name of the alarm.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The severity level can be one of emergency, alert, critical, error, warning, notice, info, and debug. This cannot be changed for some alarms

admin-enable

If disabled, the alarm is not reported in the active list and does not actuate led/failrelay.

failrelay-enable

If enabled, this alarm will assert the failrelay.

led-enable

If enabled, the main 'Alarm' LED light will be red when this alarm is asserted. If disabled, the main 'Alarm' LED light is not affected by this alarm.

6.2.3. Switch Alarm Configuration

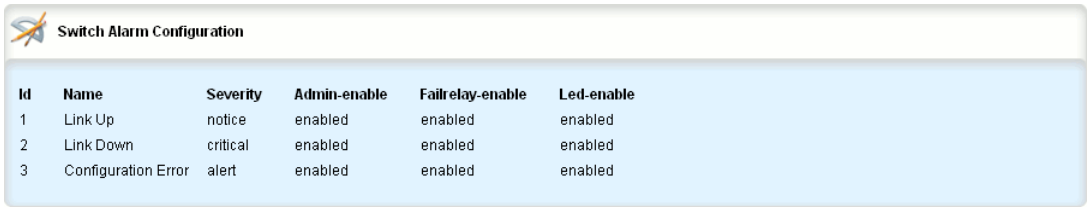

Id	Name	Severity	Admin-enable	Failrelay-enable	Led-enable
1	Link Up	notice	enabled	enabled	enabled
2	Link Down	critical	enabled	enabled	enabled
3	Configuration Error	alert	enabled	enabled	enabled

Figure 6.13. Switch Alarm Configuration table

The path to the Switch Alarm Configuration form is *admin/alarm-config/switch*.

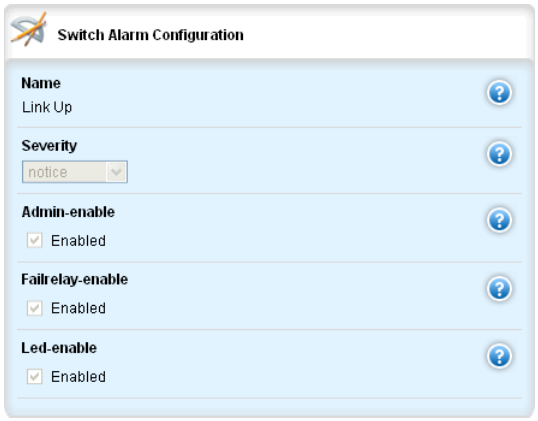

Figure 6.14. Switch Alarm Configuration form

The path to the Switch Alarm Configuration form is *admin/alarm-config/switch/{alarm id)*.

id

Synopsis: *integer*

This is the ID number of the alarm assigned by the system.

description

Synopsis: *A string*

The name of the alarm.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The severity level can be one of emergency, alert, critical, error, warning, notice, info, and debug. This cannot be changed for some alarms

admin-enable

If disabled, the alarm is not reported in the active list and does not actuate led/failrelay.

failrelay-enable

If enabled, this alarm will assert the failrelay.

led-enable

If enabled, the main 'Alarm' LED light will be red when this alarm is asserted. If disabled, the main 'Alarm' LED light is not affected by this alarm.

7. Domain Name Search

7.1. Domain Name Lookup

The DNS (Domain Name Service) menu is accessible from the main menu under *admin*. The path to this menu is *admin/dns*.

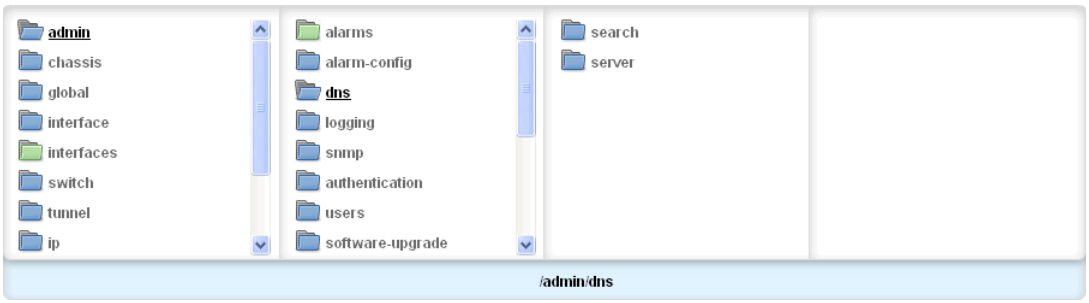

Figure 7.1. DNS menu

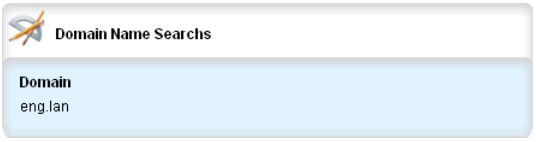

Figure 7.2. Domain Name Searches form

The path to the Domain Name Searches form is *admin/dns/search*.

domain

Synopsis: *Domain name (RFC 1034)*

Figure 7.3. Domain Name Servers

The path to the Domain Name Servers table is *admin/dns/server*.

address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

8. Logging

The syslog provides users with the ability to configure local and remote syslog connections. The remote syslog protocol, defined in RFC 3164, is a UDP/IP-based transport that enables a device to send event notification messages across IP networks to event message collectors, also known as syslog servers. The protocol is simply designed to transport these event messages from the generating device to the collector.

ROX™ supports up to 5 collectors (syslog servers). Remote Syslog provides the ability to configure:

- IP address(es) of collector(s).
- Source UDP port.
- Destination UDP port per collector.
- Syslog source facility ID per collector (same value for all ROX™ modules).
- Filtering severity level per collector (in case different collectors are interested in syslog reports with different severity levels).

8.1. Configuring Local Syslog

The local syslog configuration enables users to control what level of syslog information will be logged. Only messages of a severity level equal to or greater than the configured severity level are written to the syslog.txt file in the unit.

8.2. Configuring the Remote Syslog Server

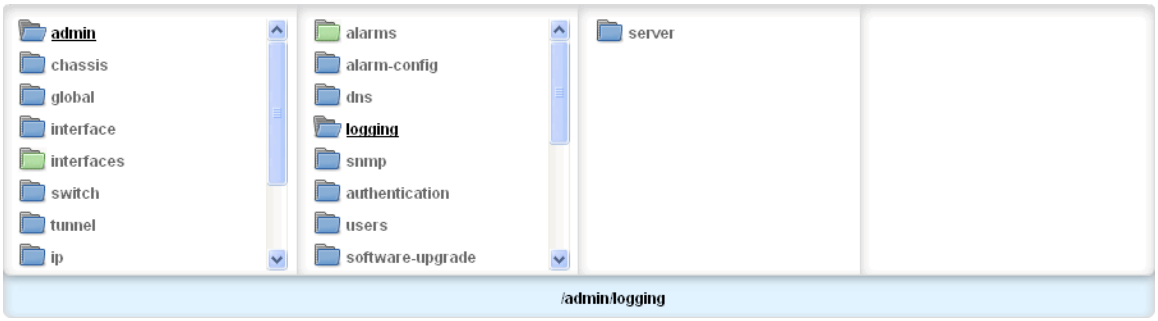

Figure 8.1. Logging menu

The Logging menu is accessible from the main menu under *admin*. The path to this menu is *admin/logging*.

Figure 8.2. Remote Server table

The Remote Server table appears on the same screen as the *Logging* menu.

The Remote Server table can be used to identify a remote logging server.

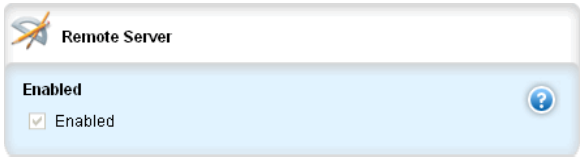

Figure 8.3. Remote Server form

If data is configured, there will be a list of logging servers under *admin/logging/server*. Clicking on each server will allow you to access the settings and Remote Server forms.

Server IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Synopsis: *Domain name (RFC 1034)*

The IPv4 or IPv6 address of a logging server. Up to 8 logging servers can be added.

enabled

Enables/disables the feed to the remote logging server

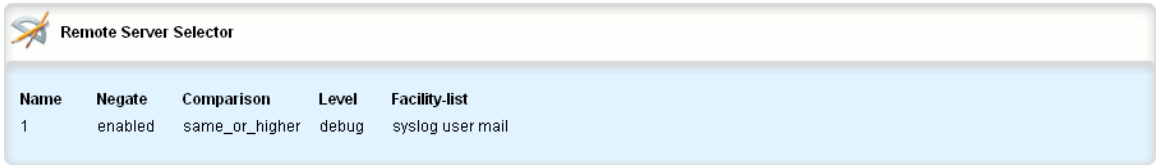

Figure 8.4. Remote Server Selector table

If data is configured, the path to the Remote Server Selector table will be *admin/logging/server*.

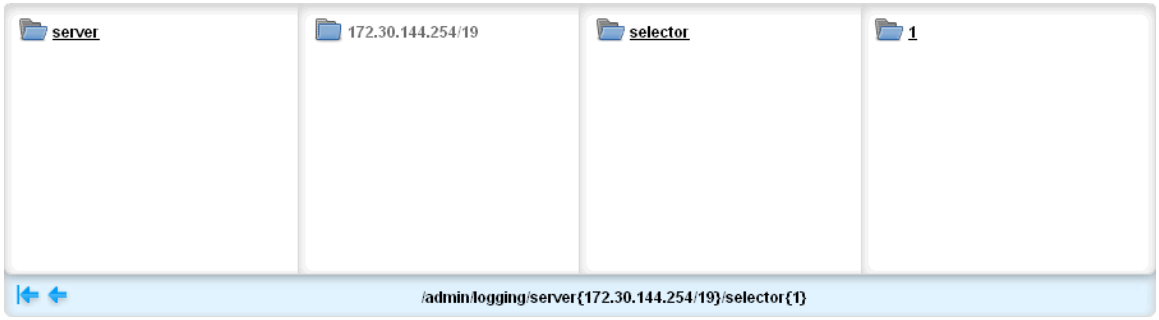

Figure 8.5. Selector menu

If data is configured, the path to the Remote Server Selector Forms (below) will be *admin/logging/server*. Then click on the next linked submenu, then on "selector" and then "1" or any linked submenus that may be in this list.

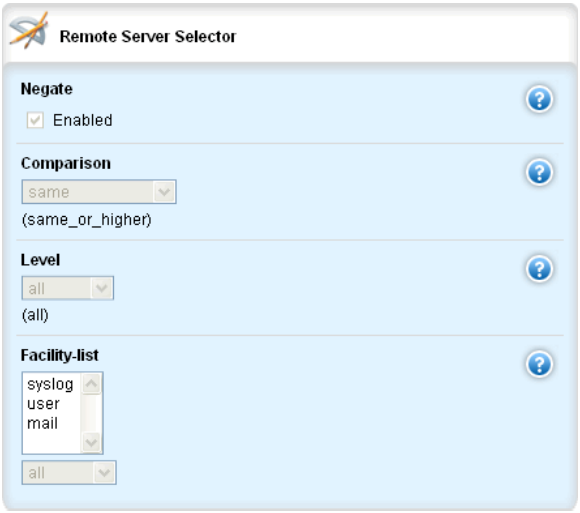

Figure 8.6. Remote Server Selector form

name

Synopsis: *integer*

The log selector identifier. Enter an integer greater than 0; up to 8 selectors can be added. The log selector determines which subsystem messages are included in the log.

negate

Excludes messages defined in the *Remote Server Selector* fields from the log. Selecting this option acts as a logical NOT for the selector definition.

For example: Selecting **same**, **debug**, and **mail** in the *Comparison*, *Level*, and *Facility-list* fields includes debug messages from the mail subsystem in the log. Selecting **Negate** *excludes* debug messages from the mail subsystem from the log.

comparison

Synopsis: *string - one of the following keywords { same, same_or_higher }*

Default: *same_or_higher*

The message severity levels to include in the log:

- **same:** includes only messages of the severity level selected in the *Level* field.
- **same_or_higher:** includes messages of the severity level selected in the *Level* field, and all messages of higher severity.

For example:

- Selecting **debug** in the *Level* field and **same** in the *Comparison* field includes only debug messages in the log.
- Selecting **debug** in the *Level* field and **same_or_higher** in the *Comparison* field includes debug and all higher severity messages in the log.

level

Synopsis: *string - one of the following keywords { all, none, debug, info, notice, warning, err, crit, alert, emerg }*

Default: *all*

The base message severity level to include in the log. **all** includes all messages. **none** excludes all messages. Other levels are listed in order of increasing severity.

facility-list

Synopsis: *string - one of the following keywords { all, local7, local6, local5, local4, local3, local2, local1, local0, uucp, user, syslog, security, news, mail, lpr, kern, ftp, daemon, cron, authpriv, auth }*

Synopsis: *"facility-list" occurs in an array of at most 8 elements.*

The subsystems generating log messages. Messages from the selected subusystems are included in the log. At least one subsystem must be selected; up to 8 subsystems can be selected.

8.3. Deleting Logs

For information on how to delete log files, see Section 2.10, “Deleting Log Files”.

9. SNMP

The SNMP (the Simple Network Management Protocol) protocol is used by network management systems and the devices they manage. SNMP is used to manage items on the device to be managed, as well as by the device itself, to report alarm conditions and other events.

The first version of SNMP, V1, provides the ability to send a notification of an event via "traps". Traps are unacknowledged UDP messages and may be lost in transit. SNMP V2 adds the ability to notify via "informs". Informs simply add acknowledgment to the trap process, resending the trap if it is not acknowledged in a timely fashion.

SNMP V1 and V2 transmit information in clear text (which may or may not be an issue depending on the facilities the data is transmitted over) and are lacking in the ability to authenticate a user. SNMP V3 adds strong authentication and encryption.

ROX™ supports Simple Network Management Protocol Version 3 (SNMPv3). This protocol provides secure access to devices by a combination of authentication and encrypting packets over the network. The security features provided are:

- message integrity - ensuring that a packet has not been tampered with in-transit.
- authentication – determining the message is from a valid source.
- encryption – scrambling the contents of a packet to prevent it from being seen by an unauthorized source.

SNMPv3 provides security models and security levels. A security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is a permitted level of security within a security model. A combination of a security model and security level will determine which security mechanism is employed when handling an SNMP packet.

Note the following about SNMPv3 protocol:

- each user belongs to a group.
- a group defines the access policy for a set of users.
- an access policy defines what SNMP objects can be accessed for: reading, writing and creating notifications.
- a group determines the list of notifications its users can receive.
- a group also defines the security model and security level for its users.

9.1. SNMP Traps

The following SNMP traps are defined in the RX1500 MIB files:

Standard	MIB	Trap and Description
RFC 3418	SNMPv2-MIB	**authenticationFailure** An authenticationFailure trap signifies that the SNMP entity has received a protocol message that is not properly authenticated. While all implementations of SNMP entities MAY be capable of generating this trap, the snmpEnableAuthenTraps object indicates whether this trap will be generated.
		coldStart A coldStart trap signifies that the SNMP entity, supporting a notification originator application, is reinitializing itself and that its configuration may have been altered.
		warmStart A warmStart trap signifies that the SNMP entity, supporting a notification originator application, is reinitializing itself such that its configuration is unaltered.

Standard	MIB	Trap and Description
RFC 4188	BRIDGE-MIB	**newRoot** The newRoot trap indicates that the sending agent has become the new root of the Spanning Tree; the trap is sent by a bridge soon after its election as the new root, e.g., upon expiration of the Topology Change Timer, immediately subsequent to its election. Implementation of this trap is optional.
		topologyChange A topologyChange trap is sent by a bridge when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Blocking state. The trap is not sent if a newRoot trap is sent for the same transition. Implementation of this trap is optional.
IEEE Std 802.1AB-2005	LLDP-MIB	**lldpRemTablesChange** An lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes. It can be utilized by an NMS to trigger LLDP remote systems table maintenance polls. Note that transmission of lldpRemTablesChange notifications are throttled by the agent, as specified by the 'lldpNotificationInterval' object.
RFC 1229, 2863, 2233, 1573	IF-MIB	**linkUp** A linkUp trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links left the down state and transitioned into some other state (but not into the notPresent state). This other state is indicated by the included value of ifOperStatus.
		linkDown A linkDown trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links is about to enter the down state from some other state (but not from the notPresent state). This other state is indicated by the included value of ifOperStatus.
RuggedCom	RUGGEDCOM-TRAPS-MIB	**trapGenericTrap** Used for "User Authentication Events" only. The main subtree for RuggedCom generic traps.
		trapPowerSupplyTrap The main subtree for RuggedCom power supply trap.
		trapSwUpgradeTrap The main subtree for RuggedCom software uppgrade trap.
		trapCfgChangeTrap The main subtree for RuggedCom configuration change trap.
		trapFanBankTrap The main subtree for RuggedCom fan bank trap.
		trapHotswapModuleStateChangeTrap The main subtree for RuggedCom fan hotswap module state change trap.

Table 9.1. SNMP Traps

9.2. SNMP Access Configuration

To configure SNMP access to ROX™, follow the procedures outlined in the example below.

9.2.1. Add an SNMP User ID

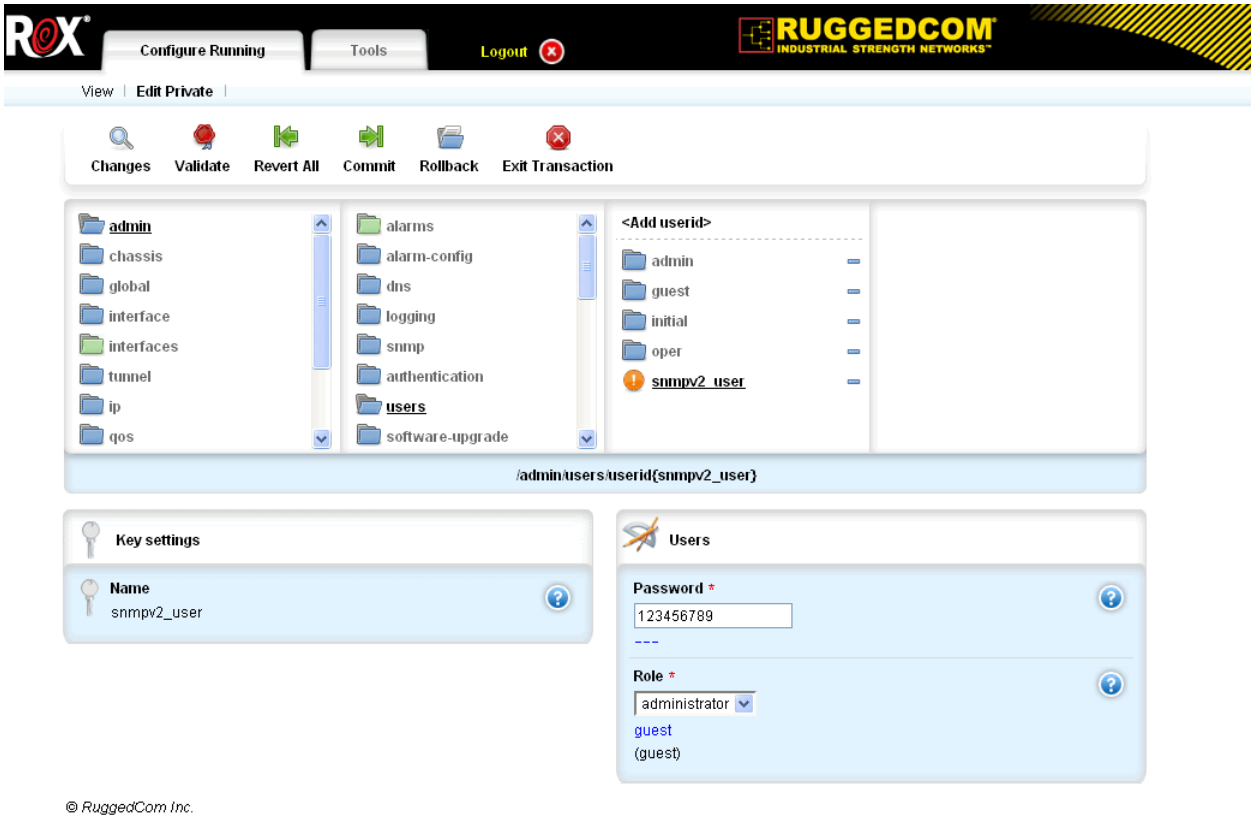

Figure 9.1. Adding an SNMP User ID

Procedure 9.1. Adding an SNMP User ID

1. Navigate to *admin/user*.
2. Click on <Add userid>. The Key settings form appears.
3. In the *Name* field, enter **snmpv2_user** and click **Add** . The Users form appears.
4. In the *Password* field, enter **123456789**.
5. In the *Role* field, enter **administrator**.
6. Click **Commit**.

9.2.2. Create an SNMP Community

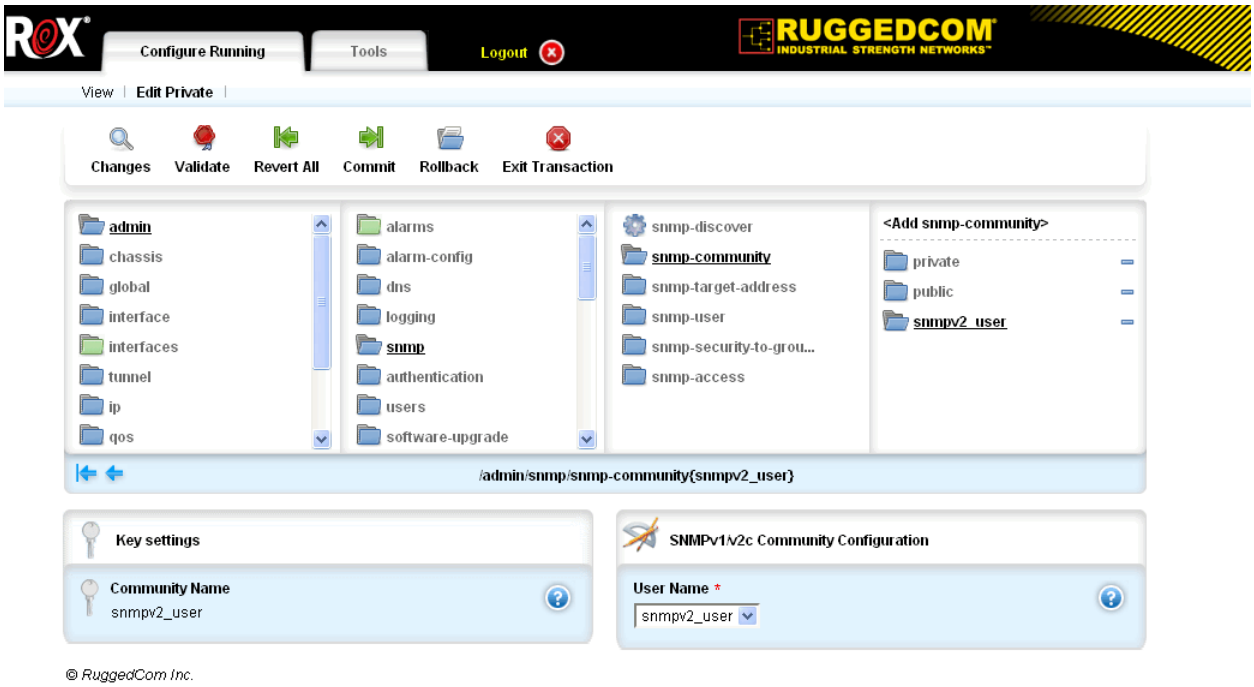

Figure 9.2. Creating an SNMP Community

Procedure 9.2. Creating an SNMP Community

1. Navigate to *admin/snmp/snmp-community*.
2. Click on <Add snmp-community>. The Key settings form appears.
3. In the *Community Name* field, enter **snmpv2_user** and click **Add**. The SNMPv1/v2c Community Configuration form appears.
4. In the *User Name* field, select **snmpv2_user**.
5. Click **Commit**.

9.2.3. Map the Community to a Security Group

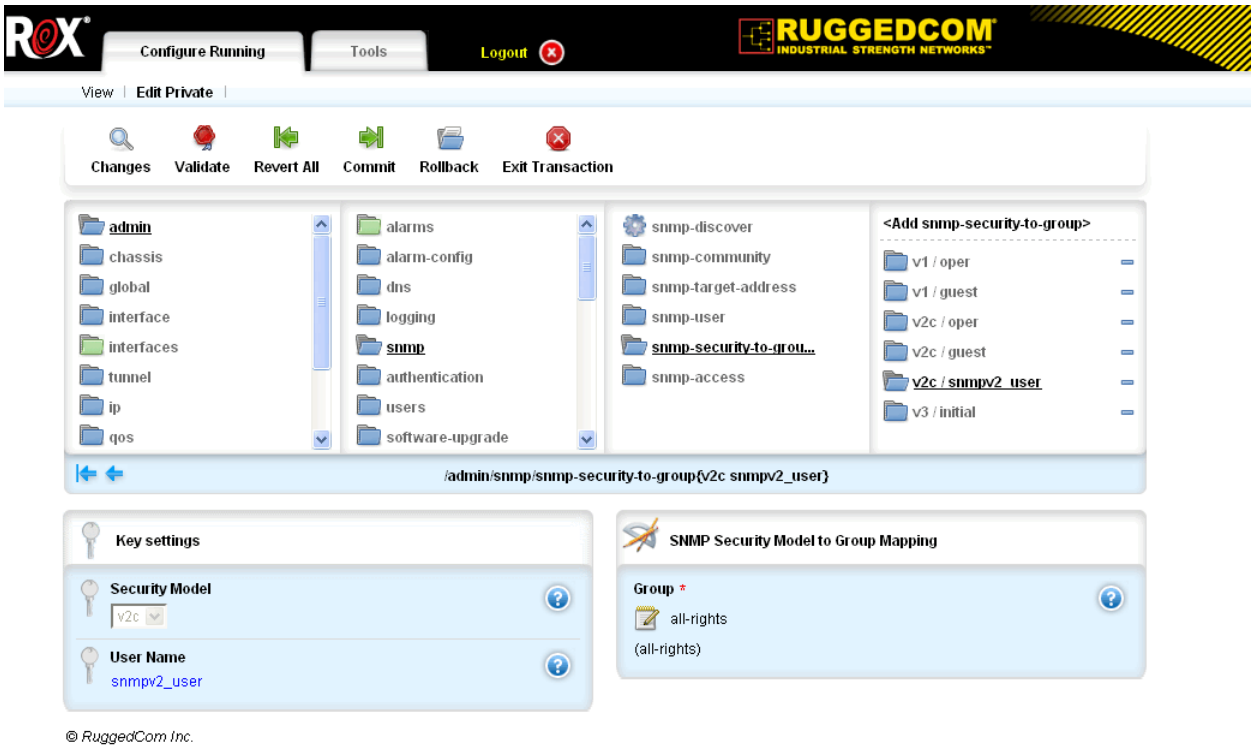

Figure 9.3. Mapping the Community to a Security Group

Procedure 9.3. Mapping the Community to a Security Group

1. Navigate to *admin/snmp/security-to-group*.
2. Click on <Add snmp-security-to-group>. The Key settings form appears.
3. In the *Security Model* field, select **v2c**.
4. In the *User Name* field, select **snmpv2_user** and click **Add**. The SNMP Security Model to Group Mapping form appears.
5. **all-rights** is the default in the Group field. Leave it as the default.
6. Click **Commit**.
7. Click **Exit Transaction**.

9.3. SNMP menu

Figure 9.4. SNMP menu

The SNMP menu is located at *admin/snmp*. The SNMP Sessions form and the SNMP USM Statistics form appear on the same screen as the *SNMP* menu.

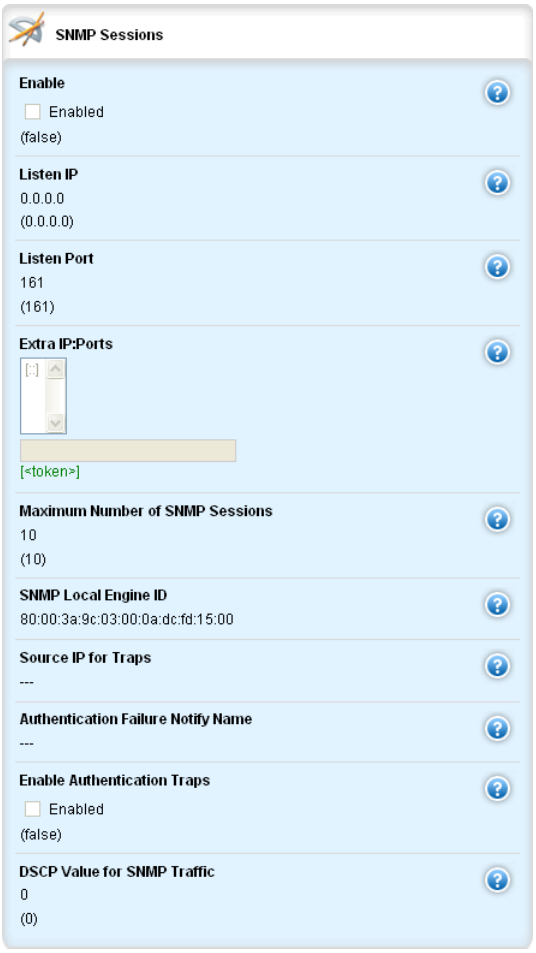

Figure 9.5. SNMP Sessions form

Enable

Synopsis: *boolean*

Default: *false*

Provides the ability to configure snmp features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the SNMP agent will listen on for SNMP requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *161*

The port the SNMP agent will listen on for SNMP requests (default 161).

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The SNMP agent will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of SNMP Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *30*

The maximum number of concurrent SNMP sessions

SNMP Local Engine ID

Synopsis: *A string*

Provides specific identification for the engine/device. By default, this value is set to use the base MAC address within the Engine ID value.

When using SNMP v3: if you change this value, you must also change the User SNMP Engine ID value for SNMP users.

Source IP for Traps

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

If set, all traffic/traps originating from this device shall use the configured IP Address for the Source IP

Authentication Failure Notify Name

Synopsis: *string - one of the following keywords { snmpv3_inform, snmpv3_trap, snmpv2_inform, snmpv2_trap, snmpv1_trap, none }*

Default: *none*

When the SNMP agent sends the standard authenticationFailure notification, it is delivered to the management targets defined for the snmpNotifyName in the snmpNotifyTable in SNMP-NOTIFICATION-MIB (RFC3413). If authenticationFailureNotifyName is the empty string (default), the notification is delivered to all management targets.

Enable Authentication Traps

Synopsis: *boolean*

Default: *false*

Enables authentication traps to be sent from the SNMP agent.

DSCP Value for SNMP Traffic

Synopsis: *unsigned byte integer in the range [0 to 63]*

Default:

Support for setting the Differentiated Services Code Point (6 bits) for traffic originating from the SNMP agent.

Figure 9.6. SNMP USM Statistics form

This table provides statistics for SNMP authentication requests

Unsupported Security Levels

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they requested a securityLevel that was unknown to the SNMP engine or otherwise unavailable.

Not In Time Windows

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they appeared outside of the authoritative SNMP engine's window.

Unknown User Names

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they referenced a user that was not known to the SNMP engine.

Unknown Engine IDs

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they referenced an snmpEngineID that was not known to the SNMP engine.

Wrong Digests

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they didn't contain the expected digest value.

Decryption Errors

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they could not be decrypted.

9.4. SNMP Discovery

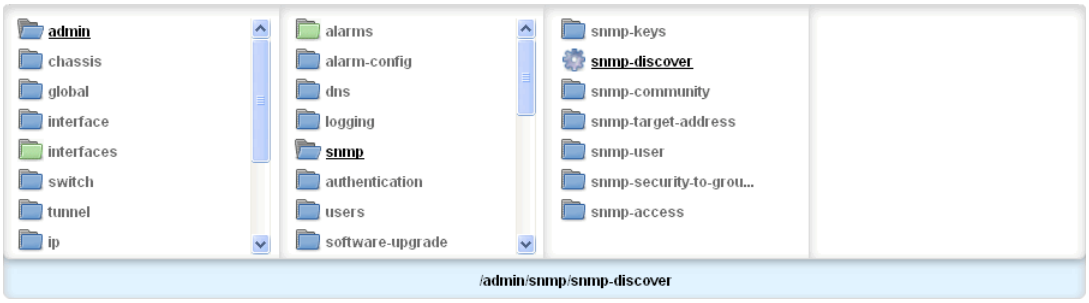

Figure 9.7. SNMP-Discover action

The path to this menu action is *admin/snmp/snmp-discover*.

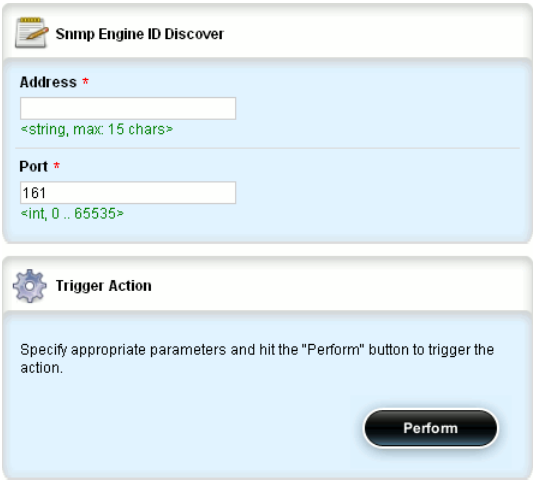

Figure 9.8. SNMP Engine ID Discover forms

To discover SNMP Engine IDs, use the SNMP Engine ID Discover and Trigger Action forms. On the *SNMP Engine ID Discover* form, enter parameters in the fields. On the *Trigger Action* form, click *Perform*.

9.5. SNMP Community

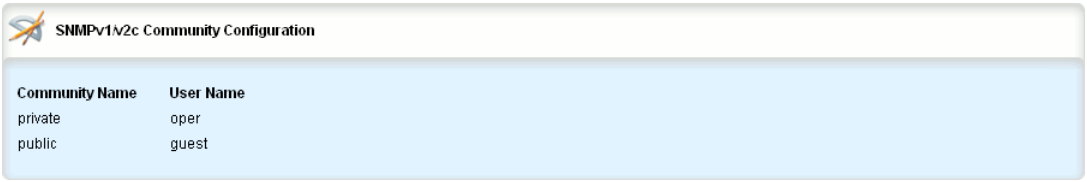

Figure 9.9. SNMPv1/v2c Community Configuration table

The path to the SNMP Community Configuration table is *admin/snmp/snmp-community*.

This table allows you to add and remove SNMPv1 and v2c Community Names.

Community Name

Synopsis: *A string*

The SNMP community name

User Name

Synopsis: *string*

The SNMP community security name

Figure 9.10. SNMPv1/v2c Community Configuration form

The path to the SNMP Community Configuration form is *admin/snmp/snmp-community/{private} or {public}*.

9.6. SNMP Target Addresses

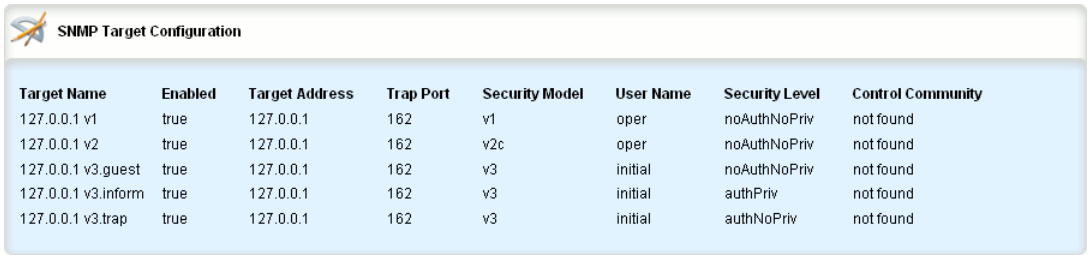

Target Name	Enabled	Target Address	Trap Port	Security Model	User Name	Security Level	Control Community
127.0.0.1 v1	true	127.0.0.1	162	v1	oper	noAuthNoPriv	not found
127.0.0.1 v2	true	127.0.0.1	162	v2c	oper	noAuthNoPriv	not found
127.0.0.1 v3.guest	true	127.0.0.1	162	v3	initial	noAuthNoPriv	not found
127.0.0.1 v3.inform	true	127.0.0.1	162	v3	initial	authPriv	not found
127.0.0.1 v3.trap	true	127.0.0.1	162	v3	initial	authNoPriv	not found

Figure 9.11. SNMP Target Configuration table

The path to the SNMP Target Configuration table is *admin/snmp/snmp-target-address*.

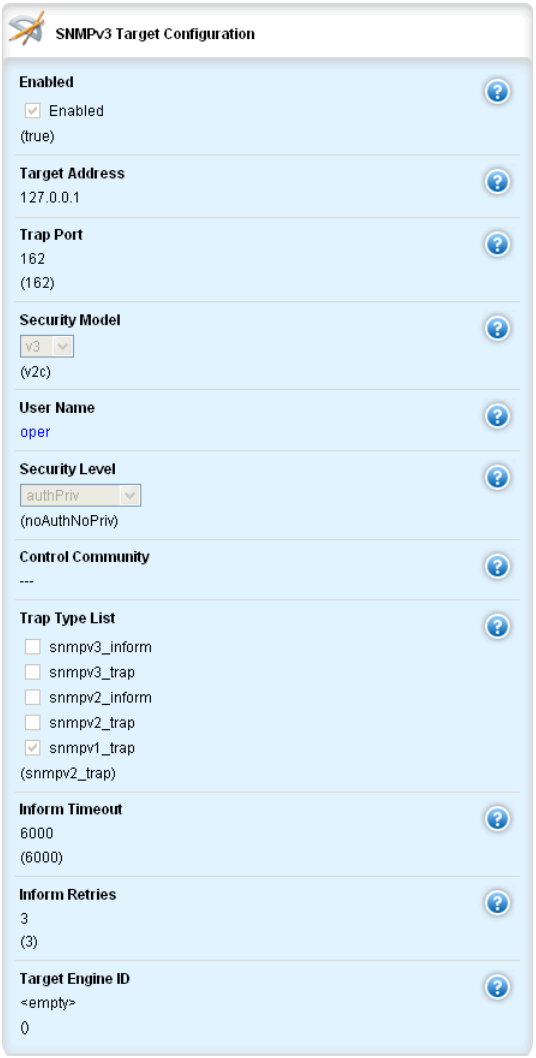

Figure 9.12. SNMPv3 Target Configuration form

To display the SNMP Target Configuration form, navigate to *admin/snmp/snmp-target-address/{address}*.

Target Name

A descriptive name for the target (ie. 'Corportate NMS')

enabled

Synopsis: *boolean*

Default: *true*

Enables/disables this specific target

Target Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

IPv4 or IPv6 address for the remote target.

Trap Port

Synopsis: *unsigned short integer*

Default: *162*

UDP Port for the remote target to receive traps on(default 162).

Security Model

Synopsis: *string - one of the following keywords { v3, v2c, v1 }*

Default: *v2c*

The SNMP security model to use: SNMPv1, SNMPv2c, or USM/SNMPv3

User Name

Synopsis: *string*

The user name to be used in communications with this target.

Security Level

Default: *noAuthNoPriv*

The SNMP security level:

- authPriv: communication with authentication and privacy.
- authNoPriv: communication with authentication and without privacy.
- noAuthnoPriv: communication without authentication and privacy.

Control Community

Synopsis: *A string*

Restrict incoming SNMP requests from the IPv4 or IPv6 address associated with this community.

Trap Type List

Default: *snmpv2_trap*

Selects the type of trap communications to be sent to this target

Inform Timeout

Default: *6000*

The timeout used for reliable inform transmissions (seconds*100)

Inform Retries

Default: *3*

The number of retries used for reliable inform transmissions

Target Engine ID

Default:

The target's SNMP local engine ID. This field may be left blank.

9.7. SNMP Users

These parameters provide the ability to configure users for the local SNMPv3 engine. Note that, if the security level employed is SNMPv1 or SNMPv2, User Name represents a community name for authentication or sending traps. Up to 32 entries can be configured.

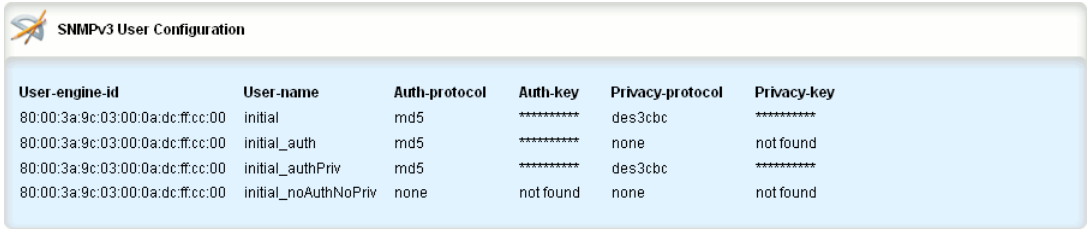

Figure 9.13. SNMP User Configuration table

The path to the SNMP User Configuration table is *admin/snmp/snmp-user*.

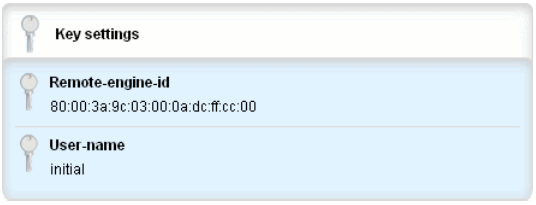

Figure 9.14. User Configuration Key Settings form

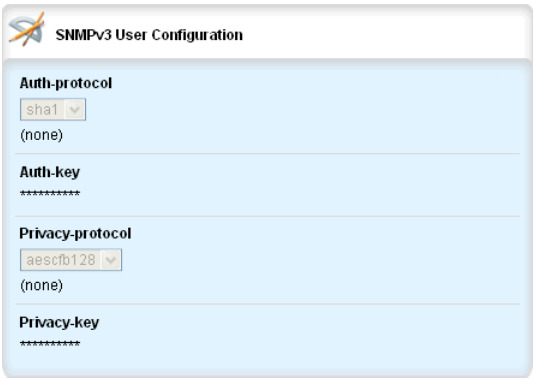

Figure 9.15. SNMP User Configuration form

The path to the Key Settings form and the SNMP User Configuration form is *admin/snmp/snmp-user/{user}*.

User SNMP Engine ID

The administratively-unique identifier for the SNMP engine; a value in the format nn:nn:nn:nn:nn:...:nn, where nn is a 2-digit hexadecimal number. Minimum length is 5 octets; maximum length is 32 octets. Each octet must be separated by a colon (:).

User Name

Synopsis: *string*

The user for the SNMP key. Select a user name from the list.

Authentication Protocol

Synopsis: *string - one of the following keywords { sha1, md5, none }*

Default: *none*

The authentication protocol providing data integrity and authentication for SNMP exchanges between the user and the SNMP engine.

Authentication Key

Synopsis: *"AES CFB128"-encrypted string*

A free-text password in the format **`$0$<your password>`**

Privacy Protocol

Synopsis: *string - one of the following keywords { aescfb128, des3cbc, none }*

Default: *none*

The symmetric privacy protocol providing data encryption and decryption for SNMP exchanges between the user and the SNMP engine.

Privacy Key

Synopsis: *"AES CFB128"-encrypted string*

A free-text password in the format **`$0$<your password>`**

9.8. SNMP Security to Group Maps

Entries in this table map the configuration of the security model and security name (user) into a group name, which is used to define an access control policy. Up to 32 entries can be configured.

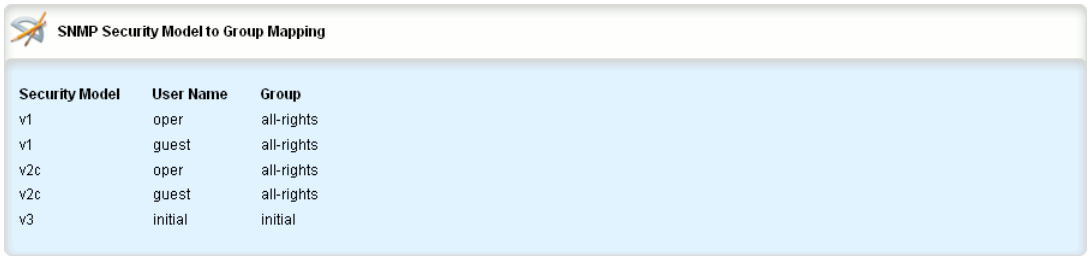

Figure 9.16. SNMP Security Model to Group Mapping table

The path to this table is *admin/snmp/snmp-security-to-group*.

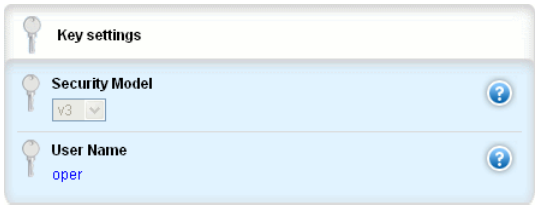

Figure 9.17. Key Settings form

Figure 9.18. SNMP Security Model to Group Mapping form

The path to these forms is *admin/snmp/snmp-security-to-group/{user}*.

Security Model

Synopsis: *string - one of the following keywords { v3, v2c, v1 }*

The SNMP security model to use: SNMPv1, SNMPv2c, or USM/SNMPv3

User Name

Synopsis: *string*

The security name (a ROX user name) for the SNMP group.

Group

Default: *all-rights*

The SNMP group name.

9.9. SNMP Access

These parameters provide the ability to configure access rights for groups. To determine whether access is allowed, one entry from this table needs to be selected and the proper view name from that entry must be used for access control checking. View names are predefined:

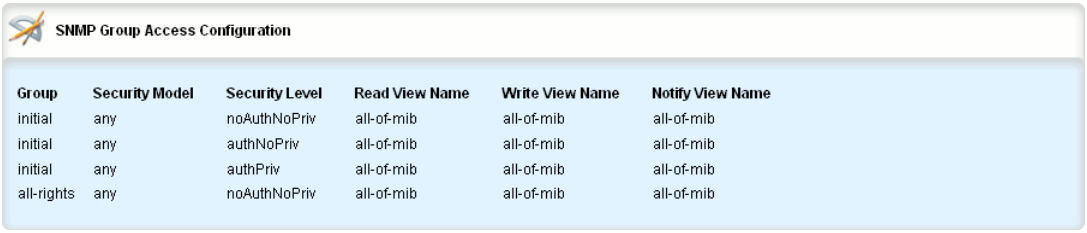
SNMP Group Access Configuration
Group Security Model Security Level Read View Name Write View Name Notify View Name
initial any noAuthNoPriv all-of-mib all-of-mib all-of-mib
initial any authNoPriv all-of-mib all-of-mib all-of-mib
initial any authPriv all-of-mib all-of-mib all-of-mib
all-rights any noAuthNoPriv all-of-mib all-of-mib all-of-mib

Figure 9.19. SNMP Group Access Configuration table

The path to this table is *admin/snmp/admin/snmp/snmp-access*.

Figure 9.20. Key Settings form

Figure 9.21. SNMP Group Access Configuration form

The path to this form is *admin/snmp/snmp-access/{access group}*.

Group

The SNMP group name.

Security Model

Synopsis: *string - one of the following keywords { v3, v2c, v1, any }*

The SNMP security model to use: SNMPv1, SNMPv2c, or USM/SNMPv3

Security Level

The SNMP security level:

- authPriv: communication with authentication and privacy.
- authNoPriv: communication with authentication and without privacy.
- noAuthnoPriv: communication without authentication and privacy.

Read View Name

Synopsis: *string - one of the following keywords { all-of-mib, restricted, v1-mib, no-view }*

Default: *all-of-mib*

The name of the read view to which the SNMP group has access: all-of-mib, restricted, v1-mib, or no-view.

Write View Name

Synopsis: *string - one of the following keywords { all-of-mib, restricted, v1-mib, no-view }*

Default: *all-of-mib*

The name of the write view to which the SNMP group has access: all-of-mib, restricted, v1-mib, or no-view.

Notify View Name

Synopsis: *string - one of the following keywords { all-of-mib, restricted, v1-mib, no-view }*

Default: *all-of-mib*

The name of the notification view to which the SNMP group has access: all-of-mib, restricted, v1-mib, or no-view.

10. Authentication

The Authentication menu is accessible from the main menu under *admin*. The path to this menu is *admin/authentication*.

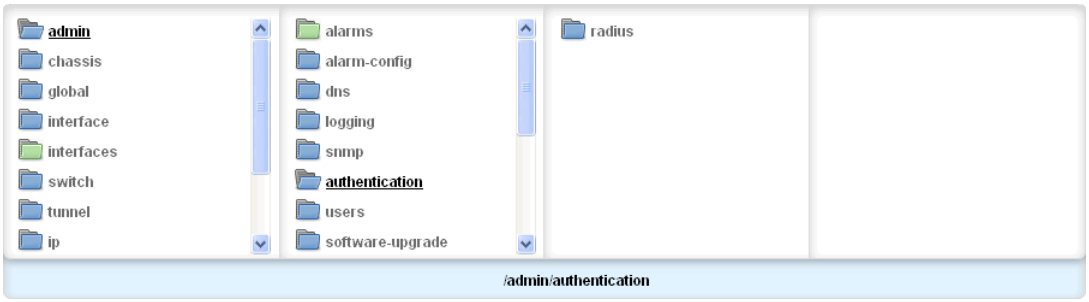

Figure 10.1. Authentication menu

The Authentication menu is accessible from the main menu under *admin*. The path to this menu is *admin/authentication*.

10.1. RADIUS

RADIUS (Remote Authentication Dial In User Service) is used to provide centralized authentication and authorization for network access. ROX™ assigns a privilege level of Admin, Operator or Guest to a user who presents a valid user name and password. The number of users who can access the ROX™ server is ordinarily dependent on the number of user records which can be configured on the server itself. ROX™ can also, however, be configured to pass along the credentials provided by the user to be remotely authenticated by a RADIUS server. In this way, a single RADIUS server can centrally store user data and provide authentication and authorization service to multiple ROX™ servers needing to authenticate connection attempts.

10.1.1. RADIUS overview

RADIUS (described in RFC 2865 [http://tools.ietf.org/html/rfc2865]) is a UDP-based protocol used for carrying authentication, authorization, and configuration information between a Network Access Server which desires to authenticate its links and a shared Authentication Server. RADIUS is also widely used in conjunction with 802.1x for port security using EAP (the Extensible Authentication Protocol, described in RFC 3748 [http://tools.ietf.org/html/rfc3748]). Refer to Chapter 24, *Port Security* for configuration details in ROX™.

A RADIUS server can act as a proxy client to other RADIUS servers or other kinds of authentication servers.

On receiving an authentication-authorization request from a client in an “Access-Request” packet, the RADIUS server checks the conditions configured for received username-password combination in the user database. If all the conditions are met, the list of configuration values for the user is placed into an “Access-Accept” packet. These values include the type of service (e.g. PPP, Login) and all the necessary values to deliver the desired service.

10.1.2. RADIUS Usage

The typical mode of operation involves a Network Access Server (NAS) - in this case the ROX™ - and a remote RADIUS server, where account information is stored. In the course of attempting to access connection-oriented services on the NAS, a user presents credentials to the NAS for authentication. The NAS forwards these to a configured RADIUS server and accepts from it the determination of whether the user is allowed the requested access. In order to protect the security of account information and of

both the NAS and the RADIUS server, transactions are encrypted and authenticated through the use of a shared secret, which is never sent in the clear.

Some administrators set the passwords of existing ROX™ accounts uniquely for each router, and then employ a common password per account for all routers served by RADIUS. The router-specific passwords are restricted to a very few personnel. A larger set of expert users is granted the rights to SSH login using the RADIUS root account passwords.

10.1.3. RADIUS on ROX™

ROX™ supports RADIUS server redundancy. Multiple RADIUS servers, usually operating from a common database, may be used to authenticate a new session. If the first configured RADIUS server does not respond, subsequent servers will be tried until a positive/negative acknowledgment is received or an attempt has been made to contact all configured servers.

Each server is configured with an associated timeout which limits the time that ROX™ will wait for a response. An authentication request could thus require up to the sum of the timeouts of all configured servers.

RADIUS authentication activity is logged to the authorization log file, “auth.log”. Details of each authentication including the time of occurrence, source and result are included.

10.1.4. RADIUS, ROX™, and Services

RADIUS provides the means to restrict access on a per-service basis. Accounts may be configured on a RADIUS server to be allowed access only to the PPP service, for example. ROX™ supports RADIUS authentication for the following services:

- *LOGIN*
- *PPP*

ROX™ provides the option of designating different servers to authenticate LOGIN or PPP services separately or in combination.

The LOGIN Service

The *LOGIN* service consists of the following types of access:

- Local console logins via the serial port and modem
- Remote shell logins via SSH and Telnet
- Secure file transfers using SCP and SFTP (based on SSH)

Authentication requests for LOGIN services will attempt to use RADIUS first. If no response is received from any configured RADIUS server, ROX™ will authenticate against the local user database.

The PPP Service

The PPP service represents incoming PPP connections via modem. Authentication requests to the PPP service use RADIUS only. In the event that no response is received from any configured RADIUS server, ROX™ will not complete the authentication request.

10.1.5. RADIUS Authentication Configuration

There are two RADIUS server forms that can be configured in ROX™.

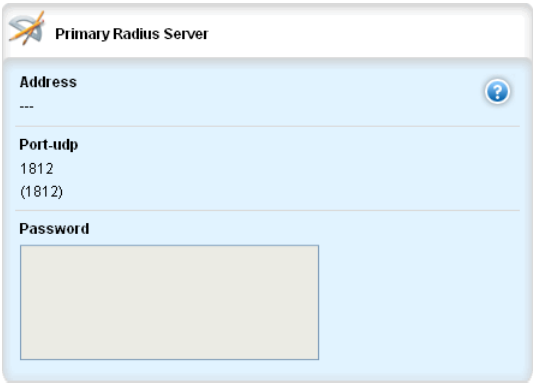

Figure 10.2. Primary RADIUS Server form

The Primary and Secondary RADIUS Server forms are accessible from the *radius* menu, which is a sub menu of the authentication menu. The path to this menu is *admin/authentication/radius*. These forms are also accessible from *global/ppp/radius*.

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

The network port of the server

password

Synopsis: *"AES CFB128"-encrypted string*

The password of the RADIUS server

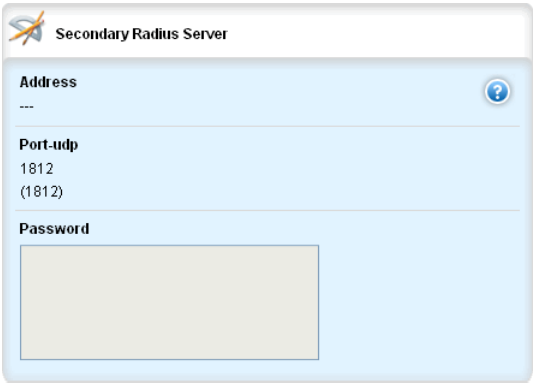

Figure 10.3. Secondary RADIUS Server form

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

The network port of the server

password

Synopsis: *"AES CFB128"-encrypted string*

The password of the RADIUS server

For more information on 802.1x Authentication, please see Chapter 24, *Port Security*. For additional information on RADIUS server configuration, please see Appendix B, *RADIUS Server Configuration*.

11. NETCONF

Figure 11.1. NETCONF menu

The NETCONF menu is accessible from the main menu under *admin*. The path to this menu is *admin/netconf*.

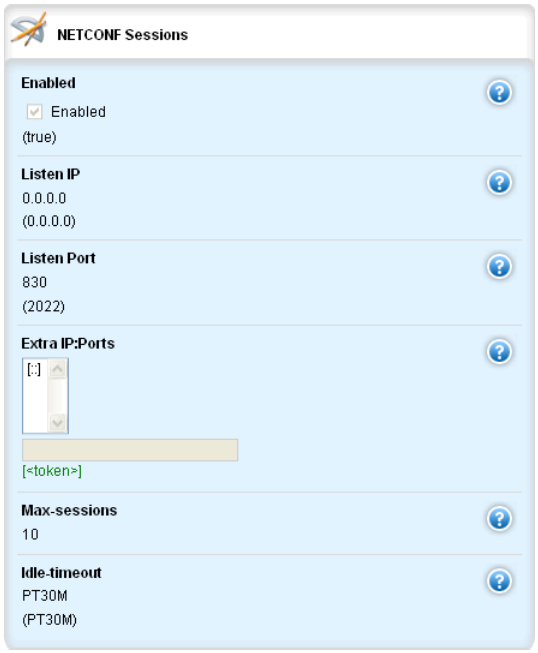

Figure 11.2. NETCONF Sessions form

The path to the NETCONF Sessions form and the NETCONF State/Statistics form is *admin/netconf*.

enabled

Synopsis: *boolean*

Default: *true*

Provides the ability to configure NETCONF features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the CLI will listen on for NETCONF requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *830*

The port on which NETCONF listens for NETCONF requests. The default is port 830.

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

Additional IP addresses and ports on which NETCONF listens for NETCONF requests. You can specify IP addresses and ports in the following forms:

- nnn.nnn.nnn.nnn:port represents an IPv4 address followed by a colon and port number. For example, 192.168.10.12:19343
- [::] represents the default IPv4 address and default port number. This is the default configuration.
- [::]:port represents an IPv6 address followed by a colon and port number. For example, [fe80::5eff:35ff]:16000

Maximum Number of NETCONF Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *10*

The maximum number of concurrent NETCONF sessions

Idle Timeout

Default: *PT0S*

Maximum idle time before terminating a NETCONF session. If the session is waiting for notifications, or has a pending confirmed commit, the idle timeout is not used. The default value is 0, which means no timeout.

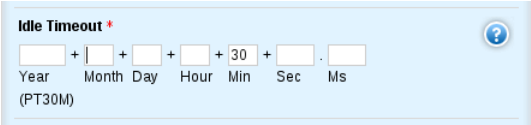

Figure 11.3. Idle-timeout field

Clicking on the Idle-timeout field on the NETCONF Sessions form allows you to choose a value for this field. The default value is PT30M, which stands for "Precision Time 30 Minutes". This refers to the time when an inactive session expires or times out. Only integer values corresponding to the following fields can be entered: Year, Month, Day, Hour, Min, Sec, or Ms. The example above shows the default value of PT30M, which corresponds to the Min field.

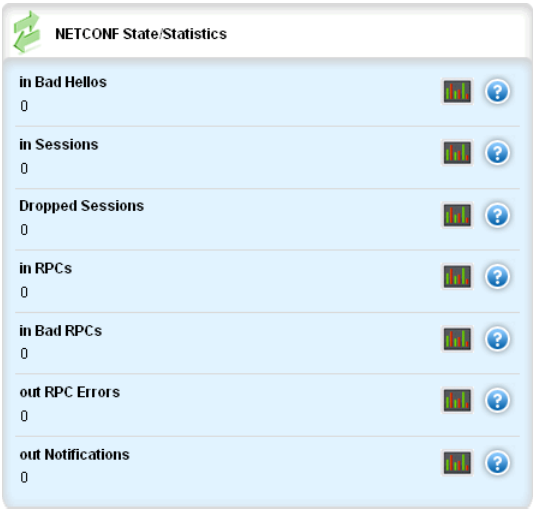

Figure 11.4. NETCONF State/Statistics form

in Bad Hellos

Synopsis: *unsigned integer*

The total number of sessions silently dropped because an invalid 'hello' message was received. This includes hello messages with a 'session-id' attribute, bad namespace, and bad capability declarations.

in Sessions

Synopsis: *unsigned integer*

The total number of NETCONF sessions started towards the NETCONF peer.

inSessions - inBadHellos = 'number of correctly started netconf sessions'

Dropped Sessions

Synopsis: *unsigned integer*

The total number of NETCONF sessions dropped.

inSessions - inBadHellos = 'number of correctly started netconf sessions'

in RPCs

Synopsis: *unsigned integer*

The total number of rpc requests received.

in Bad RPCs

Synopsis: *unsigned integer*

The total number of rpcs which were parsed correctly, but couldn't be serviced because they contained non-conformant XML.

out RPC Errors

Synopsis: *unsigned integer*

The total number of 'rpc-reply' messages with 'rpc-error' sent.

out Notifications

Synopsis: *unsigned integer*

The total number of 'notification' messages sent.

12. Chassis Management

This chapter describes basic system commands to control the ROX™-based system. Among the data available are:

- Hardware type and revision information
- Module firmware versions
- Current module status
- Temperature and power supply sensor data

This chapter also describes the “Chassis” menu and sub menus, which provide comprehensive diagnostic information for the RuggedBackbone™ chassis and all installed hardware modules.

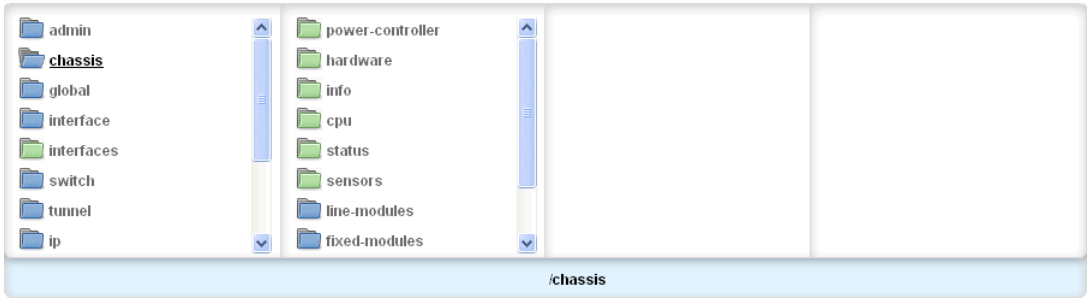

Figure 12.1. Chassis menu

The Chassis menu is accessible from the main menu under *Chassis*. This menu contains chassis-level configuration and status features. A variety of sub menus can be linked to from the Chassis menu. The Chassis sub menu section is organized so that information tables appear on a screen closer to the top level and clicking on one of the sub menus brings up the form(s) associated with the table. For example, clicking on the Chassis menu and then the Hardware sub menu will display the Slot Hardware table. Further clicking on the pm1 sub menu will display the Slot Hardware form.

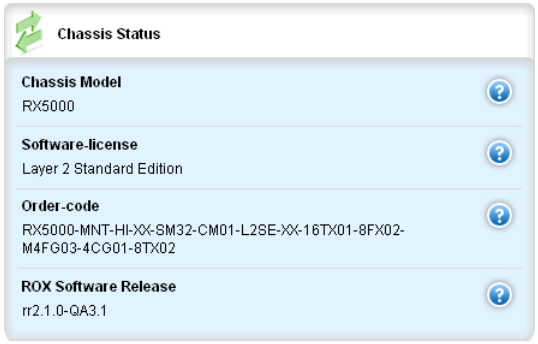

Figure 12.2. Chassis Status form

The Chassis Status form contains basic status information about the chassis. This form appears on the same screen as the *Chassis* menu.

Chassis Model

Synopsis: *string*

The RuggedCom device model name.

software-license

Synopsis: *string*

The current software capability.

order-code

Synopsis: *A string*

The order code derived from the current configuration of the device.

ROX Software Release

Synopsis: *string*

The release of ROX running on the chassis.

12.1. Power Controller

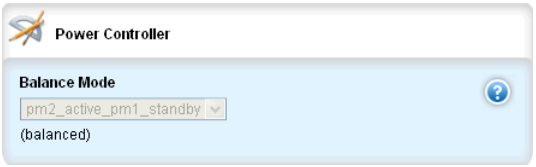

Figure 12.3. Power Controller form

As of ROX version 2.2, the balance-mode feature is not supported. This feature remains in the interface for backwards compatibility.

balance-mode

synopsis: string - one of the following keywords { PM2_Active_PM1_Standby, PM1_Active_PM2_Standby, Balanced }

When more than one power modules are present, this parameter specifies how they share the provision of power to the chassis. Select the "Balanced" option to cause each PM to provide an equal amount of power, Select the "PM1_Active_PM2_Standby" or "PM2_Active_PM1_Standby" options to provide the power from the active PM. The default value of this parameter is "Balanced"

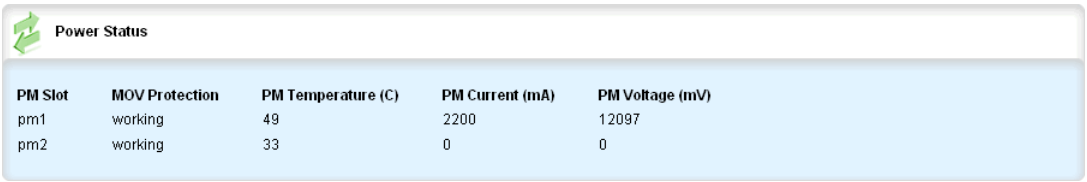

Figure 12.4. Power Status table

This table displays status information for power modules.

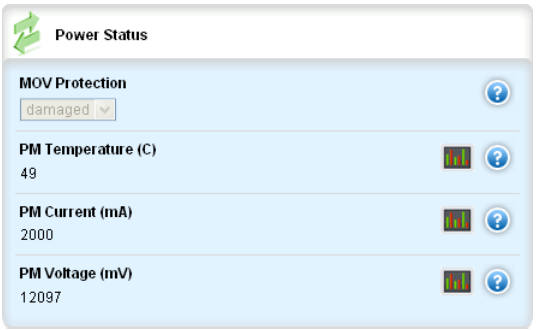

Figure 12.5. Power Status form

PM Slot

Synopsis: *string - one of the following keywords { pm2, pm1 }*

The name of the power module slot as labeled on the chassis

MOV Protection

Synopsis: *string - one of the following keywords { damaged, working, na }*

The state of the MOV protection circuit

PM Temperature (C)

Synopsis: *integer*

The temperature (Celsius) inside the power module

PM Current (mA)

Synopsis: *integer*

The current (mA) sourced by the power module

PM Voltage (mV)

Synopsis: *integer*

The voltage (mV) sourced by the power module

12.2. Slot Hardware

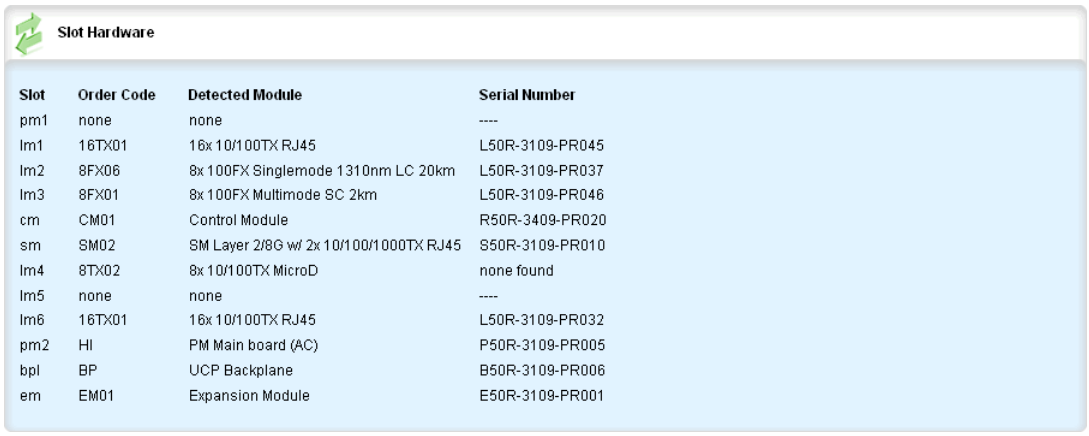

Slot Hardware

Slot	Order Code	Detected Module	Serial Number
pm1	none	none	----
lm1	16TX01	16x 10/100TX RJ45	L50R-3109-PR045
lm2	8FX06	8x 100FX Singlemode 1310nm LC 20km	L50R-3109-PR037
lm3	8FX01	8x 100FX Multimode SC 2km	L50R-3109-PR046
cm	CM01	Control Module	R50R-3409-PR020
sm	SM02	SM Layer 2/8G w/ 2x 10/100/1000TX RJ45	S50R-3109-PR010
lm4	8TX02	8x 10/100TX MicroD	none found
lm5	none	none	----
lm6	16TX01	16x 10/100TX RJ45	L50R-3109-PR032
pm2	HI	PM Main board (AC)	P50R-3109-PR005
bpl	BP	UCP Backplane	B50R-3109-PR006
em	EM01	Expansion Module	E50R-3109-PR001

Figure 12.6. Slot Hardware table

Figure 12.7. Slot Hardware form

The Slot Hardware table and form contain identifying information about the hardware module installed in a particular chassis slot.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Order Code

Synopsis: *A string*

The order code of the chassis as derived from the current hardware configuration.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

Serial Number

Synopsis: *A string*

The installed module's unique serial number.

12.3. Slot Identification

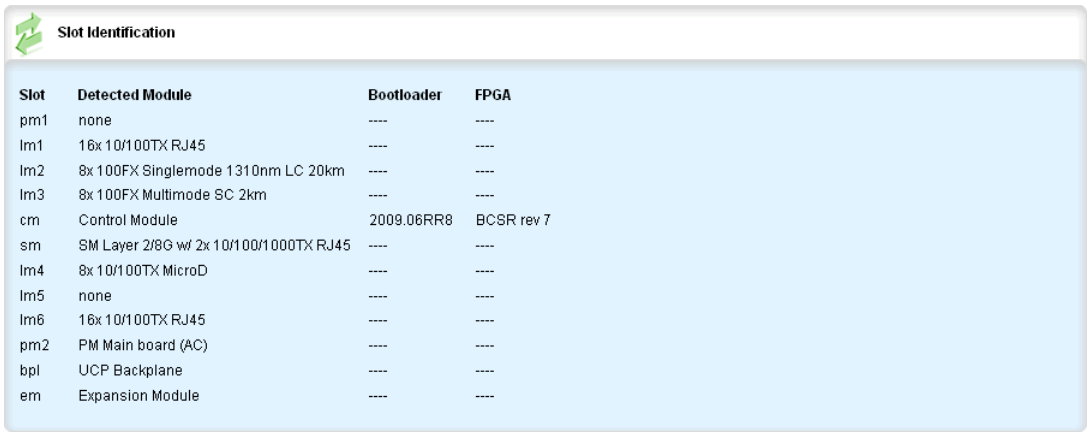
Slot Identification
Slot Detected Module Bootloader FPGA
pm1 none ---- ----
lm1 16x 10/100TX RJ45 ---- ----
lm2 8x 100FX Singlemode 1310nm LC 20km ---- ----
lm3 8x 100FX Multimode SC 2km ---- ----
cm Control Module 2009.06RR8 BCSR rev 7
sm SM Layer 2/8G w/ 2x 10/100/1000TX RJ45 ---- ----
lm4 8x 10/100TX MicroD ---- ----
lm5 none ---- ----
lm6 16x 10/100TX RJ45 ---- ----
pm2 PM Main board (AC) ---- ----
bpl UCP Backplane ---- ----
em Expansion Module ---- ----

Figure 12.8. Slot Identification table

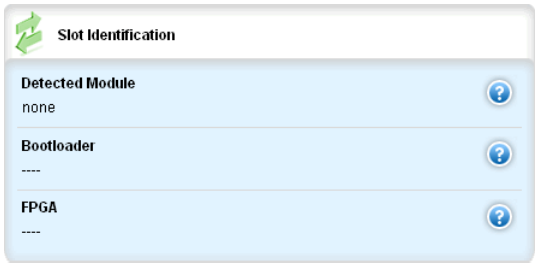

Figure 12.9. Slot Identification form

The Slot Identification table and form contain version information about the module software installed in a particular chassis slot.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

Bootloader

Synopsis: *string*

The version of the ROX bootloader software on the installed module.

FPGA

Synopsis: *string*

The version of the ROX FPGA firmware (if any) running on the installed module.

12.4. CPU

This section deals with CPU and memory usage for each slot (not all modules have a cpu). The Slot CPU table and form display status information about the hardware module installed in a particular chassis slot. The path to the Slot CPU/RAM Utilization table is *chassis/cpu*.

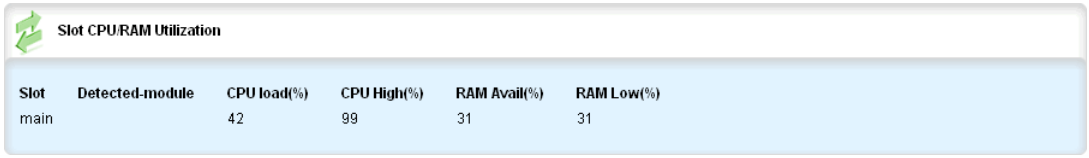

Figure 12.10. Slot CPU/RAM Utilization table

The path to the Slot CPU/RAM Utilization form is *chassis/cpu* and then clicking on a linked sub menu (for example, lm1).

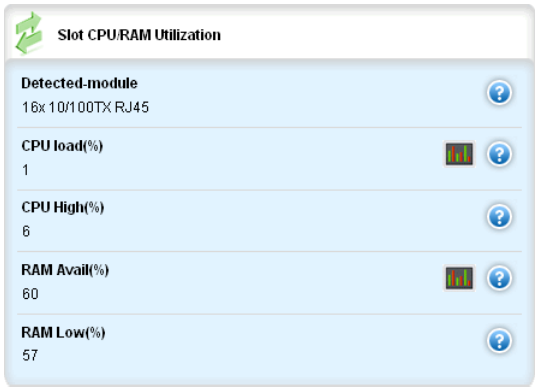

Figure 12.11. Slot CPU/RAM Utilization form

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

detected-module

Synopsis: *A string*

The installed module's type specifier.

CPU load(%)

Synopsis: *integer*

The CPU load, in percent, on the installed module.

RAM Avail(%)

Synopsis: *integer*

The proportion of memory (RAM) currently unused, in percent, on the installed module.

RAM Low(%)

Synopsis: *integer*

The lowest proportion of unused memory (RAM), in percent, recorded for the installed module since start-up.

12.5. Slot Status

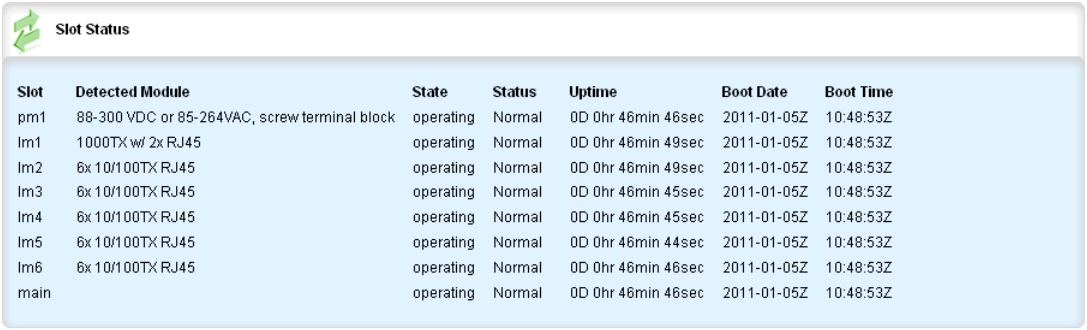
Slot Status
Slot Detected Module State Status Uptime Boot Date Boot Time
pm1 88-300 VDC or 85-264VAC, screw terminal block operating Normal 0D 0hr 46min 46sec 2011-01-05Z 10:48:53Z
lm1 1000TX w/ 2x RJ45 operating Normal 0D 0hr 46min 49sec 2011-01-05Z 10:48:53Z
lm2 6x 10/100TX RJ45 operating Normal 0D 0hr 46min 49sec 2011-01-05Z 10:48:53Z
lm3 6x 10/100TX RJ45 operating Normal 0D 0hr 46min 45sec 2011-01-05Z 10:48:53Z
lm4 6x 10/100TX RJ45 operating Normal 0D 0hr 46min 45sec 2011-01-05Z 10:48:53Z
lm5 6x 10/100TX RJ45 operating Normal 0D 0hr 46min 44sec 2011-01-05Z 10:48:53Z
lm6 6x 10/100TX RJ45 operating Normal 0D 0hr 46min 46sec 2011-01-05Z 10:48:53Z
main operating Normal 0D 0hr 46min 46sec 2011-01-05Z 10:48:53Z

Figure 12.12. Slot Status table

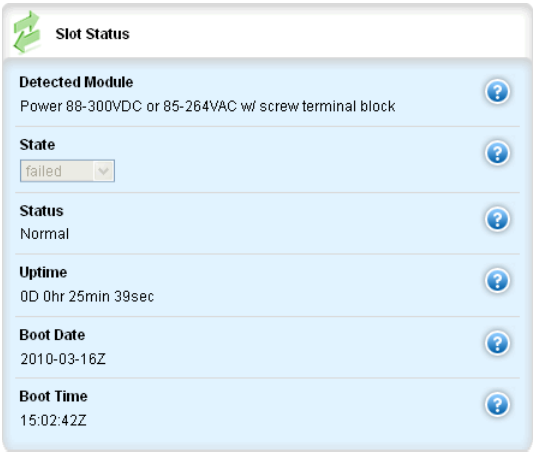

Figure 12.13. Slot Status form

The Slot Status table and form display status information about the hardware module installed in a particular chassis slot.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

State

Synopsis: *string - one of the following keywords { disconnected, failed, operating, resetting, disabled, empty, unknown }*

The current state of the installed module.

Status

Synopsis: *string*

The runtime status of the installed module.

Uptime

Synopsis: *string*

The total time elapsed since the start-up of the installed module.

Boot Date

Synopsis: *date specification*

The date on which the installed module was started up.

Boot Time

Synopsis: *time specification*

The time at which the installed module was started up.

12.6. Slot Sensors

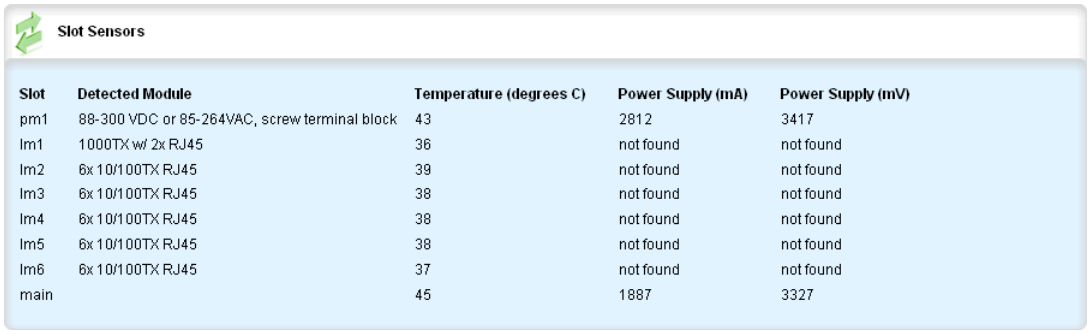

Figure 12.14. Slot Sensors table

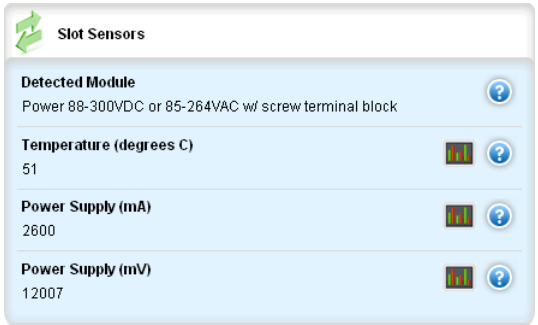

Figure 12.15. Slot Sensors form

Slot sensors contain temperature and power supply information about the installed modules.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

Temperature (degrees C)

Synopsis: *integer*

The temperature, in degrees C, of the installed module. If multiple temperature sensors are present on the board, the maximum reading is reported.

Power Supply (mA)

Synopsis: *integer*

The power supply current, in mA, being drawn by the installed module.

Power Supply (mV)

Synopsis: *integer*

The power supply voltage, in mV, seen by the installed module.

12.7. Module Configuration

For information on adding and replacing line modules, see Appendix D, Adding and Replacing Line Modules.

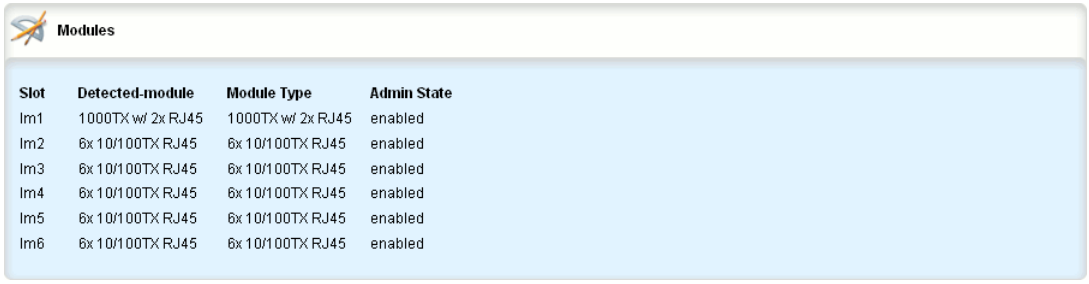
Modules
Slot Detected-module Module Type Admin State
lm1 1000TX w/ 2x RJ45 1000TX w/ 2x RJ45 enabled
lm2 6x 10/100TX RJ45 6x 10/100TX RJ45 enabled
lm3 6x 10/100TX RJ45 6x 10/100TX RJ45 enabled
lm4 6x 10/100TX RJ45 6x 10/100TX RJ45 enabled
lm5 6x 10/100TX RJ45 6x 10/100TX RJ45 enabled
lm6 6x 10/100TX RJ45 6x 10/100TX RJ45 enabled

Figure 12.16. Modules table

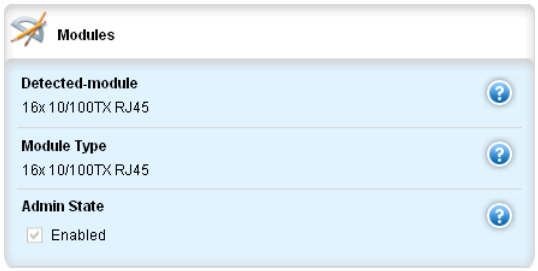

Figure 12.17. Modules form

The Module Configuration feature provides administrative control of the installed modules. The Modules table and form provide information about the administrative control of a module in a particular chassis slot.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The slot name, as marked on the silkscreen across the top of the chassis.

detected-module

Synopsis: *A string*

The installed module's type specifier.

Module Type

Synopsis: *A string*

Sets the module type to be used in this slot.

Admin State

Sets the administrative state for a module. Enabling the module powers it on.

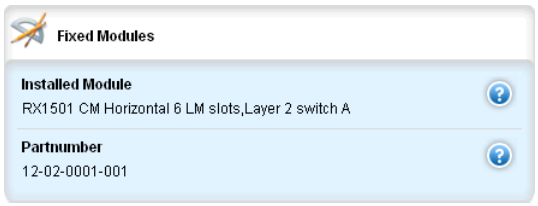

Figure 12.18. Fixed Modules form

Figure 12.19. Fixed Modules table

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Installed Module

Synopsis: *A string*

The module type to be used in this slot.

partnumber

Synopsis: *A string*

The part number of the module type in this slot.

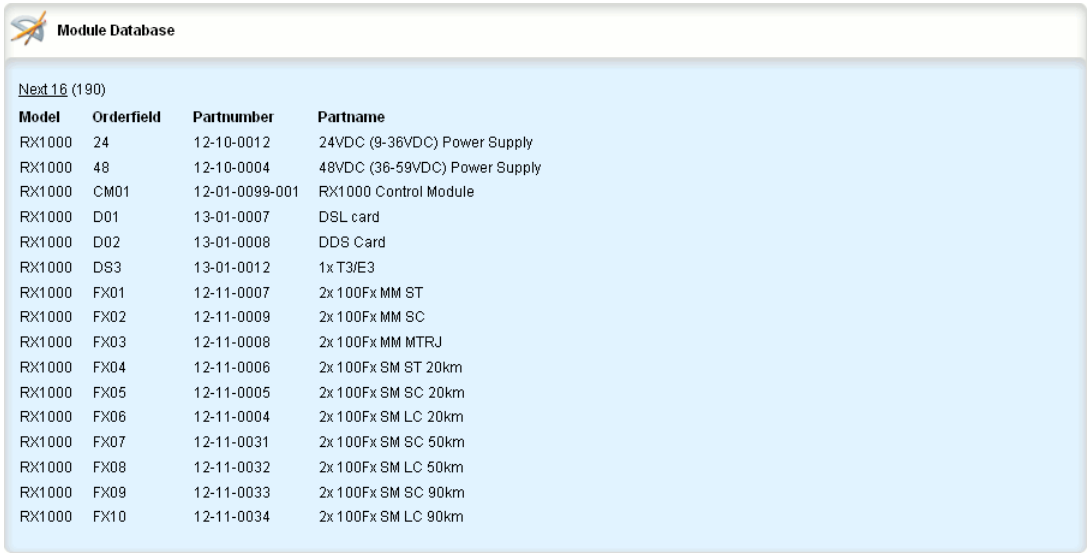

Module Database

Next 16 (190)

Model	Orderfield	Partnumber	Partname
RX1000	24	12-10-0012	24VDC (9-36VDC) Power Supply
RX1000	48	12-10-0004	48VDC (36-59VDC) Power Supply
RX1000	CM01	12-01-0099-001	RX1000 Control Module
RX1000	D01	13-01-0007	DSL card
RX1000	D02	13-01-0008	DDS Card
RX1000	DS3	13-01-0012	1x T3/E3
RX1000	FX01	12-11-0007	2x 100Fx MM ST
RX1000	FX02	12-11-0009	2x 100Fx MM SC
RX1000	FX03	12-11-0008	2x 100Fx MM MTRJ
RX1000	FX04	12-11-0006	2x 100Fx SM ST 20km
RX1000	FX05	12-11-0005	2x 100Fx SM SC 20km
RX1000	FX06	12-11-0004	2x 100Fx SM LC 20km
RX1000	FX07	12-11-0031	2x 100Fx SM SC 50km
RX1000	FX08	12-11-0032	2x 100Fx SM LC 50km
RX1000	FX09	12-11-0033	2x 100Fx SM SC 90km
RX1000	FX10	12-11-0034	2x 100Fx SM LC 90km

Figure 12.20. Module Database table

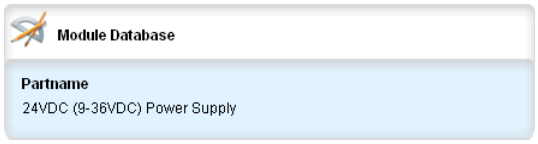

Module Database

Partname

24VDC (9-36VDC) Power Supply

Figure 12.21. Module Database form

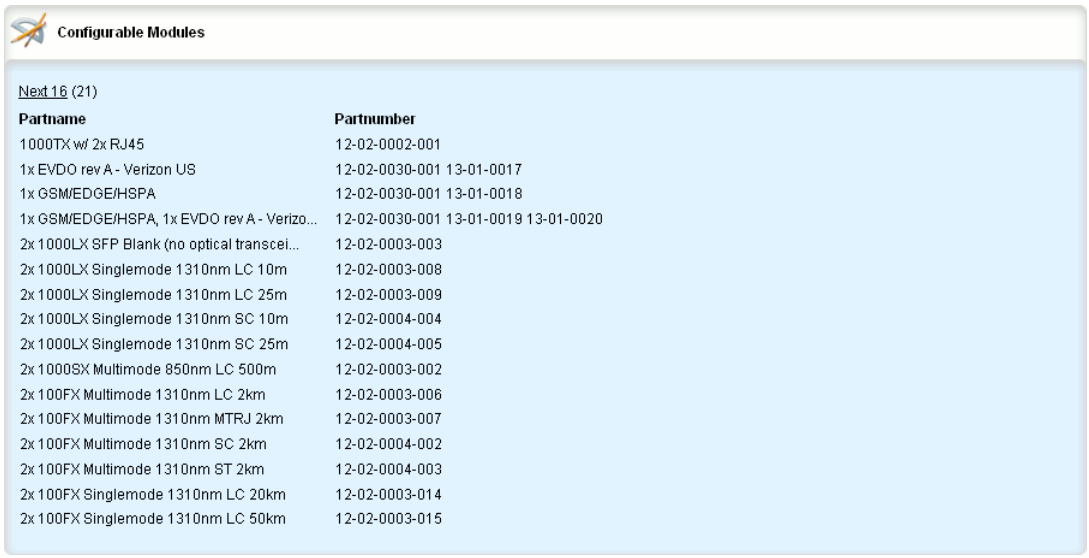

Configurable Modules

Next 16 (21)

Partname	Partnumber
1000TX w/ 2x RJ45	12-02-0002-001
1x EVDO rev A - Verizon US	12-02-0030-001 13-01-0017
1x GSM/EDGE/HSPA	12-02-0030-001 13-01-0018
1x GSM/EDGE/HSPA, 1x EVDO rev A - Verizo...	12-02-0030-001 13-01-0019 13-01-0020
2x 1000LX SFP Blank (no optical transcei...	12-02-0003-003
2x 1000LX Singlemode 1310nm LC 10m	12-02-0003-008
2x 1000LX Singlemode 1310nm LC 25m	12-02-0003-009
2x 1000LX Singlemode 1310nm SC 10m	12-02-0004-004
2x 1000LX Singlemode 1310nm SC 25m	12-02-0004-005
2x 1000SX Multimode 850nm LC 500m	12-02-0003-002
2x 100FX Multimode 1310nm LC 2km	12-02-0003-006
2x 100FX Multimode 1310nm MTRJ 2km	12-02-0003-007
2x 100FX Multimode 1310nm SC 2km	12-02-0004-002
2x 100FX Multimode 1310nm ST 2km	12-02-0004-003
2x 100FX Singlemode 1310nm LC 20km	12-02-0003-014
2x 100FX Singlemode 1310nm LC 50km	12-02-0003-015

Figure 12.22. Configurable Modules table

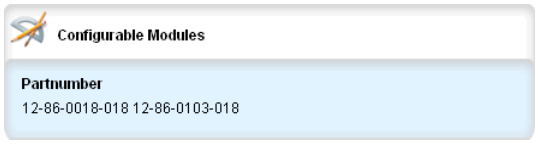

Configurable Modules

Partnumber

12-86-0018-018 12-86-0103-018

Figure 12.23. Configurable Modules form

13. PPP Users

13.1. Overview

Use the PPP menu to configure local and remote authentication for PPP user login through an L2TP client.

A PPP Server can be configured to accept a connection request only after validating the user's credentials. When so configured, the login credentials can be verified locally based on the information configured in *global/ppp/profiles/dial-in*, or can be verified through a remote RADIUS server configured in *global/ppp/radius*.

Use the dial-out option to create a profile for users who dial out to a PPP server. In this case, the device acts as a PPP client.

PPP users, profiles, and settings are configured on forms found under the PPP menu. To display the PPP menu, navigate to *global/ppp*.

13.2. PPP Configuration

The PPP menu is accessible from the main menu under *admin*. The PPP menu provides access to forms that allow you to add and remove PPP users and profiles and make adjustments to related settings.

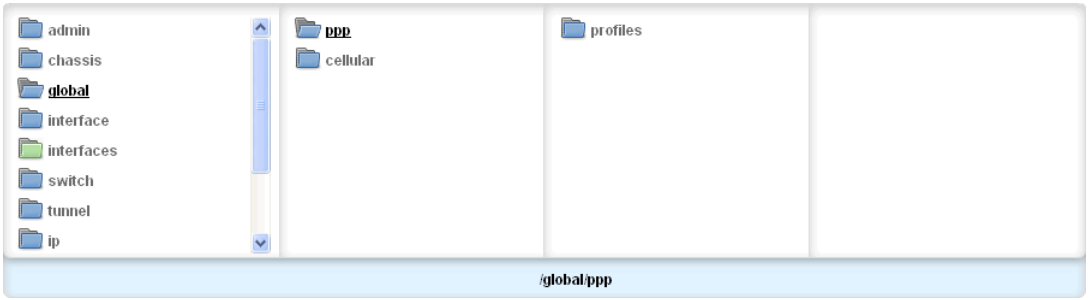

Figure 13.1. PPP menu

To display the Dial-in PPP Users table, navigate to *global/ppp/profiles/dialin*.

Figure 13.2. Dial-in PPP Users table

name

Synopsis: *A string*

The user ID of the remote client used to be authenticated

password

Synopsis: *A string*

The password of the remote client used to be authenticated

To display the Dial-in Users form, navigate to *global/ppp/profiles/dialin*, followed by any of the linked submenus.

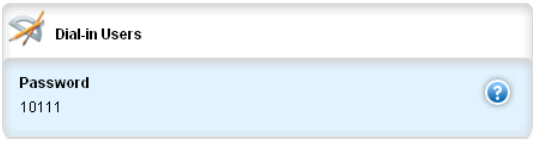

Figure 13.3. Dial-in Users form

The Dial-in Users form allows you to add PPP profiles for dial-in users.

To display the Dial-out PPP Users table, navigate to *global/ppp/profiles/dialout*.

Figure 13.4. Dial-out PPP Users table

Dial-out PPP is used to add PPP profile for dialOut users.

name

Synopsis: *A string*

The connection name

To display the PPP Configuration form, navigate to *global/ppp/profiles/dialout* and then click on any of the linked submenus.

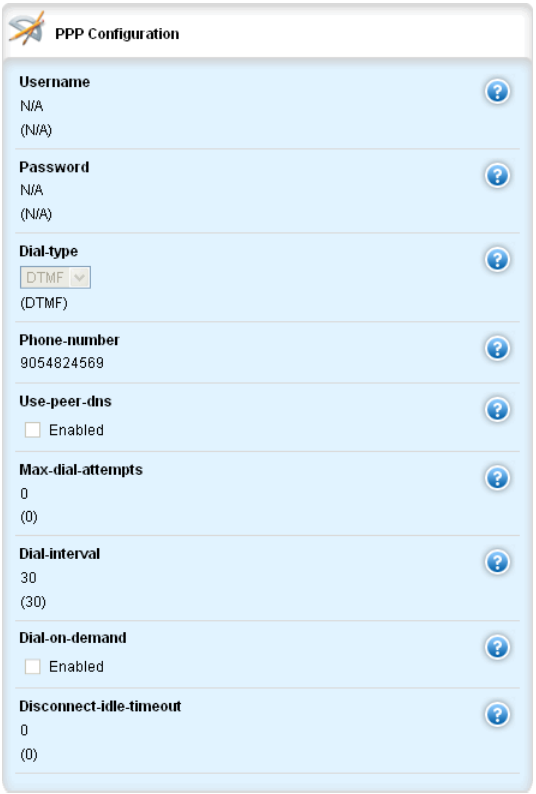

Figure 13.5. PPP Configuration form

username

Synopsis: *A string*

Default: *N/A*

The user ID used to log on to a remote PPP server

password

Synopsis: *A string*

Default: *N/A*

The password used to log on to a remote PPP server

dial-type

Synopsis: *string - one of the following keywords { Pulse, DTMF }*

Default: *DTMF*

The type of dialing system to use on the phone line. This field is only applied on analog modems

phone-number

Synopsis: *A string*

The telephone number to dial to connect to the remote PPP server

use-peer-dns

Using DNS recommended by the PPP server

max-dial-attempts

Synopsis: *integer*

Default:

The number of consecutive times that the modem will dial the phone number before it stops attempting to establish a connection. Zero (0) means the modem will try to connect to the PPP server forever.

dial-interval

Synopsis: *integer*

Default: *30*

The time in seconds to wait before re-initiating the link after it terminates

dial-on-demand

Activate Dial-on-Demand on this connection. The establishment of the PPP connection is postponed until there is data to be tranmitted via the interface

disconnect-idle-timeout

Synopsis: *integer*

Default:

The time in seconds to wait before disconnecting PPP when there is no traffic on the link. This option is only valid when Dial-on-Demand is enabled

failover-on-demand

Activate link failover on-demand on this device. PPP establishment on this device is controlled by link failover

To display the PPP RADIUS configuration forms, navigate to *global/ppp/radius*.

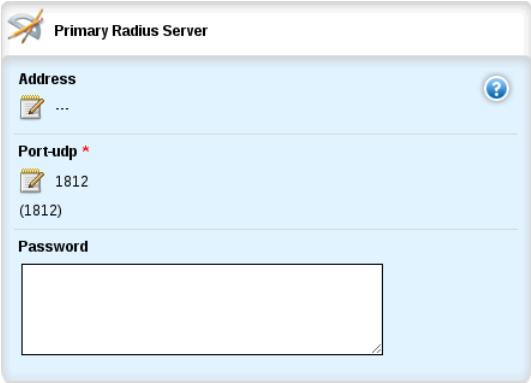

Figure 13.6. PPP Primary Radius Server form

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

password

Synopsis: *"AES CFB128"-encrypted string*

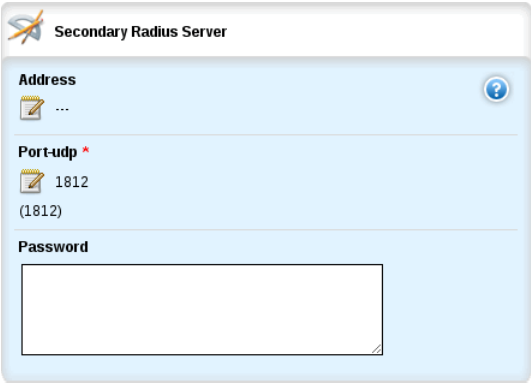

Figure 13.7. PPP Secondary Radius Server form

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

password

Synopsis: *"AES CFB128"-encrypted string*

13.3. PPP Interfaces and Link Failover

Link failover provides an easily configured means of raising a backup link upon the failure of a designated main link. A PPP interface can be specified as the backup link for a designate main link. For more information on link failover, see Chapter 41, *Link Failover*.

The PPP **Dial-on-demand** option is a standard PPP option. This option triggers the modem dial-out when there is traffic passing through the modem link. The modem hangs up when traffic stops within the time set in the PPP **Disconnect-idle-timeout** option. When **Dial-on-demand** is enabled, the presence of traffic controls the operation of the modem link.

When using a PPP interface with link failover, the link failover **On-demand** option allows link failover to bring up or take down the PPP interface as needed. Link failover triggers the modem dial-out to establish a PPP connection and pass traffic over the modem link when the main link is down. When the main link is up again, link failover takes down the PPP interface. The PPP **Disconnect-idle timeout** setting does not apply to the PPP interface when the interface is brought up by link failover.

*The link failover **On-demand** option and the PPP **Dial-on-demand** option are mutually exclusive: only one of these options should be set for a PPP interface.*

14. DHCP Relay

A DHCP Relay Agent is a device that forwards DHCP packets between clients and servers when they are not on the same physical LAN segment or IP subnet. The feature is enabled if the DHCP server IP address and a set of access ports are configured.

DHCP Option 82 provides a mechanism for assigning an IP Address based on the location of the client device in the network. Information about the client's location can be sent along with the DHCP request to the server. Based on this information, the DHCP server makes a decision about an IP Address to be assigned.

DHCP Relay Agent takes the broadcast DHCP requests from clients received on the configured access port and inserts the relay agent information option (Option 82) into the packet. Option 82 contains the VLAN ID (2 bytes) and the port number of the access port (2 bytes: the circuit ID sub-option) and the switch's MAC address (the remote ID sub-option). This information uniquely defines the access port's position in the network. For example, in ROX™, the Circuit ID for VLAN 2 on LM 4 Port 15 is **`00:02:04:0F`**.

The DHCP Server supporting DHCP Option 82 sends a unicast reply and echoes Option 82. The DHCP Relay Agent removes the Option 82 field and broadcasts the packet to the port from which the original request was received.

The DHCP Relay Agent communicates to the server on a management interface. The agent's IP address is the address configured for the management interface.

ROX™ can be configured to act as a DHCP Relay Agent that forwards DHCP and BOOTP requests from clients on one layer 2 network to one or more configured DHCP servers on other networks. This allows the implementation of some measure of isolation between DHCP clients and servers.

The DHCP Relay Agent is configured to listen for DHCP and BOOTP requests on particular Ethernet and VLAN network interfaces, and to relay to a list of one or more DHCP servers. When a request is received from a client, ROX™ forwards the request to each of the configured DHCP servers. When a reply is received from a server, ROX™ forwards the reply back to the originating client.

While DHCP Relay and DHCP Server may both be configured to run concurrently, they may not be configured to run on the same network interface.

Figure 14.1. DHCP Relay Agent Menu

The DHCP Relay Agent menu is available under *switch/dhcp-relay-agent*.

Figure 14.2. DHCP Relay Agent Form

The DHCP Relay Agent form appears on the same screen as the *DHCP Relay Agent* menu.

DHCP Server Address

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of the DHCP server to which DHCP queries will be forwarded from this relay agent.

Figure 14.3. DHCP Relay Agent Client Ports table

To display the DHCP Relay Agent Client Ports table, navigate to *dhcp-relay-agent/dhcp-client-ports*.

DHCP Relay Agent Client Ports are ports where DHCP clients are connected.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

15. DHCP Server

15.1. DHCP Fundamentals

Dynamic Host Configuration Protocol (DHCP) is a method for centrally and consistently managing IP addresses and settings for clients, offering a variety of assignment methods. IP addresses can be assigned based on the Ethernet MAC address of a client, sequentially, or by using port identification provided by a DHCP relay agent device.

15.1.1. DHCP Network Organizations

The information to assign addresses in DHCP is organized to deal with clients at the interface, subnet, pool, shared network, host-group, and host levels.

Interfaces specify the IP interface to which the client sends a request.

Subnets control settings for each subnet that DHCP serves. A subnet can include a range of IP address to hand out to clients. Subnets contain groups, pools and hosts. Only one subnet can contain dynamic IP address ranges without any access restrictions on any given physical port, since DHCP doesn't know which subnet a client should belong to when the request is received.

Pools contain ranges of IP addresses to hand out to clients with access rules to determine which clients should receive addresses from that pool.

Shared networks are used when multiple subnets should be served by a single physical port. This applies both when using a DHCP relay agent connected to the port with additional subnets behind the relay agent, or when multiple virtual networks exist on one physical interface. Each subnet then gets its own subnet definition inside the shared network rather than at the top level. Shared networks contain subnets, groups and hosts.

Host-groups allow identical settings to be created for a group of hosts, making it easier to manage changes to the settings for all the hosts contained within the group. Groups contain hosts.

Host entries assign settings to a specific client based on its Ethernet MAC address.

15.1.2. Option 82 Support with Disable NAK

If DHCP relay clients (option 82 clients) are used on the same subnet as the DHCP server, some clients will try to renew a lease immediately after receiving it by requesting a renewal directly from the DHCP server. Because the DHCP server is only configured to provide the lease through a relay agent configured with the correct Option 82 fields, the server sends the client an NAK to disallow the lease. Enabling Option 82 disables the reject message so that the renewal request sent from the DHCP relay agent (which the DHCP server accepts since it has the correct Option 82 fields added) is the only message for which the client receives a reply. If the DHCP server and clients are not on the same subnet, this option is not required.

Note that the meaning of the value of many fields depends on the client's interpretation of the field, so the actual meaning of a field is determined by the client. To determine which values are required by the client for special options, refer to the client documentation.

15.2. Configuring DHCP Server

The DHCP Server menu is available under *services* at *services/dhcpserver*.

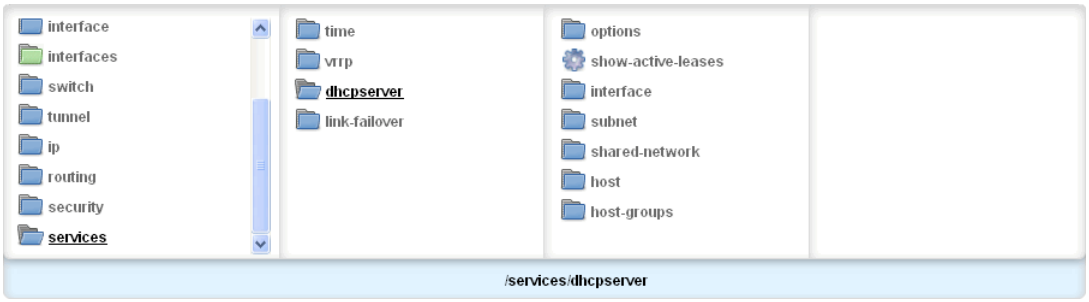

Figure 15.1. DHCP Server menu

Under *services/dhcpserver*, you can configure the following:

- enable the DHCP service. See Section 15.2.1, “Enabling the DHCP Service”.
- set the interface. See Section 15.2.2, “DHCP Interfaces”.
- configure subnets and pools. See Section 15.2.3, “DHCP Subnets and Pools”.
- configure shared networks. See Section 15.2.4, “DHCP Shared Networks”.
- configure hosts. See Section 15.2.5, “DHCP Hosts”.
- configure host-groups. See Section 15.2.6, “DHCP Host-groups”.
- configure DHCP options. See Section 15.2.8, “DHCP Options”.

Under *services/dhcpserver*, you can also view a list of active DHCP leases. See Section 15.2.7, “Viewing Active DHCP Leases”.

15.2.1. Enabling the DHCP Service

To enable or disable the DHCP service, navigate to *services/dhcpserver*. On the *Dynamic Host Control Protocol (DHCP) server* form, enable or disable the DHCP service.

Figure 15.2. DHCP Server form

enabled

Enables and disables the the DHCP server.

15.2.2. DHCP Interfaces

A *DHCP listen interface* is the interface on which DHCP listens for lease requests. You can configure multiple DHCP interfaces.

- To view a list of DHCP listen interfaces, navigate to *services/dhcpserver/interface*.

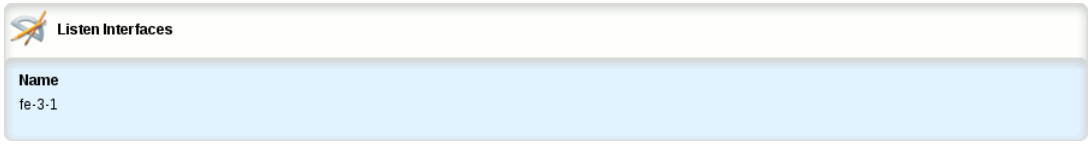

Figure 15.3. Listen Interfaces table

- To add a DHCP listen interface, enter edit mode, navigate to *services/dhcpserver/interface*, and click **<Add interface>**. On the *Key settings* form, select an interface from the list and click **Add**.

15.2.3. DHCP Subnets and Pools

- To view a list of DHCP subnets, navigate to */services/dhcpserver/subnet*.

Figure 15.4. Subnet Configuration table

- To add a subnet, enter edit mode, navigate to */services/dhcpserver/subnet*, and click **<Add subnet>**. On the *Key settings* form, type a name for the subnet and click **Add**.

  On the *Subnet Configuration* form, set the subnet parameters.

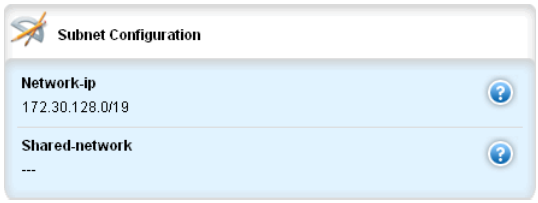

Figure 15.5. Subnet Configuration form

network-ip

Synopsis: *IPv4 address and prefix in CIDR notation*

The network IP address for this subnet

shared-network

Synopsis: *A string*

The shared-network that this subnet belongs to

You can configure DHCP options at the subnet level. Options set at this level override options set at higher levels.

- To set Lease Configuration and Client Configuration options, navigate to */services/dhcpserver/subnet{subnet id}/options*. For more information, see Section 15.2.8.1, “Lease Configuration Options” and Section 15.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/subnet{subnet id}/options/client*. For more information, see Section 15.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/subnet{subnet id}/options/client/custom* and click <Add custom>. For more information, see Section 15.2.9, “Custom DHCP Options”.

15.2.3.1. DHCP Pools

- To view a list of DHCP pools, navigate to */services/dhcpserver/subnet{subnet02}/options/iprange*.

Figure 15.6. IP Pool Configuration table

- To add a DHCP pool, enter edit mode, navigate to */services/dhcpserver/subnet{subnet02}/options/iprange*, and click <Add iprange>. On the *Key settings* form, type the starting IP address of the range and click **Add**. On the *IP pool configuration* form, type the ending IP address of the range.

15.2.4. DHCP Shared Networks

- To view a list of shared networks, navigate to *services/dhcpserver/shared-network*.

Figure 15.7. Shared Network Configuration table

- To add a shared network, enter edit mode, navigate to *services/dhcpserver/shared-network*, and click **<Add shared-network>**. On the *Key settings* form, set a name for the shared network and click **Add**.

You can configure DHCP options at the shared network level. Options set at this level override options set at higher levels.

- To set Lease Configuration and Client Configuration options, navigate to */services/dhcpserver/shared-network{shared network id}/options*. For more information, see Section 15.2.8.1, “Lease Configuration Options” and Section 15.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/shared-network{shared network id}/options/client*. For more information, see Section 15.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/shared-network{shared network id}/options/client/custom* and click <Add custom>. For more information, see Section 15.2.9, “Custom DHCP Options”.

15.2.5. DHCP Hosts

- To view a list of DHCP hosts, navigate to *services/dhcpserver/host*.

Figure 15.8. Host Configuration table

- To add a DHCP host, enter edit mode, navigate to *services/dhcpserver/host*, and click **<Add host>**. On the *Key settings* form, type a name for the host and click **Add**.

You can configure DHCP options at the host level. Options set at this level override options set at higher levels.

- To set Hardware Configuration, Lease Configuration, and Client Configuration options, navigate to */services/dhcpserver/host{host id}/options*. For more information, see Section 15.2.10, “Hardware Configuration”, Section 15.2.8.1, “Lease Configuration Options”, and Section 15.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/host{host id}/options/client*. For more information, see Section 15.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/host{host id}/options/client/custom* and click <Add custom>. For more information, see Section 15.2.9, “Custom DHCP Options”.

15.2.6. DHCP Host-groups

- To view a list of host-groups, navigate to *services/dhcpserver/host-groups*.

Figure 15.9. Host Group Configuration table

- To add a host-group, enter edit mode, navigate to */services/dhcpserver/host-groups*, and click **<Add host-groups>**. On the *Key settings* form, type a name for the host-group and click **Add**.

You can configure DHCP options at the host-group level. Options set at this level override options set at higher levels.

- To set Lease Configuration and Client Configuration options, navigate to */services/dhcpserver/host-groups{host-group id}/options* . For more information, see Section 15.2.8.1, “Lease Configuration Options” and Section 15.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/host-groups{host-group id}/options/client*. For more information, see Section 15.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/host-groups{host-group id}/options/client/custom* and click <Add custom>. For more information, see Section 15.2.9, “Custom DHCP Options”.

15.2.7. Viewing Active DHCP Leases

You can view a list of active DHCP leases. The list includes the start and end times, hardware ethernet address, and client hostname for each lease.

To view DHCP leases, navigate to *services/dhcpserver* and click **show-active-leases**. On the *Trigger Action* form, click **Perform**. The */services/dhcpserver/show-active-leases* form appears and displays the active DHCP leases.

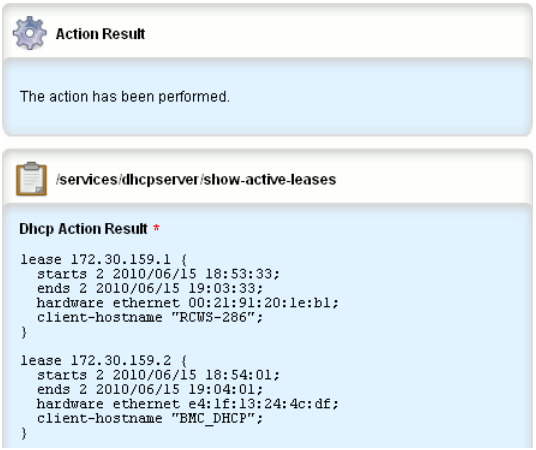

Figure 15.10. /services/dhcpserver/show-active-leases form

15.2.8. DHCP Options

You can set DHCP options at the subnet, shared network, host-groups, and hosts level. Options set at lower levels override those set at higher levels.

DHCP options are set on the following forms:

- *Leased Configuration* form: sets the DHCP lease times. This form is used at all DHCP levels.
- *Client Configuration* form at the DHCP grouping level: sets client address, host-group, and other settings. Subnets and shared networks use the same form; hosts and host-groups have unique settings on this form.
- *Client Configuration* form at the DHCP option level: sets hostname, subnet, DNS, and other settings. This form is used at all DHCP levels.
- *NIS Configuration* form: sets NIS server information. This form is used at all DHCP levels.
- *NetBios Configuration* form: sets NetBios scope and nameserver information. This form is used at all DHCP levels.
- *Custom* form: sets DHCP custom options. This form is used at all DHCP levels.
- *Hardware Configuration*: sets network type and MAC address information. This form is used for hosts only.

15.2.8.1. Lease Configuration Options

You can set the lease configuration options at all DHCP levels. To set DHCP lease configuration options, enter edit mode and navigate to:

- DHCP server options: */services/dhcpserver/options*
- subnet options: */services/dhcpserver/subnet{subnet id}/options*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options*
- host-group options: */services/dhcpserver/host-groups{host-group id}/options*
- host options: */services/dhcpserver/host{host id}/options*

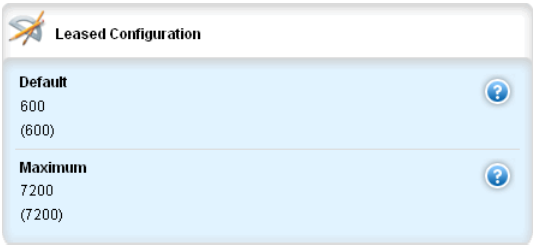

Figure 15.11. Lease Configuration form

default

Synopsis: *integer*

Default: *600*

The minimum leased time that the server offers to the client

maximum

Synopsis: *integer*

Default: *7200*

The maximum leased time that the server offers to the client

15.2.8.2. Client Configuration Options at the DHCP Levels

Different DHCP options are set at the subnet and shared network, hosts, and host groups levels.

15.2.8.2.1. Client Configuration Options: Subnets and Shared Networks

To set DHCP client configuration options at the subnet and shared networks levels, enter edit mode and navigate to:

- subnet options: */services/dhcpserver/subnet{subnet id}/options*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options*

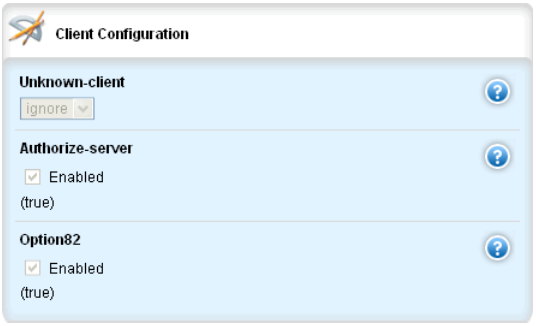

Figure 15.12. Client Configuration form for Subnets and Shared Networks

unknown-client

Synopsis: *string - one of the following keywords { ignore, deny, allow }*

The action to take for previously unregistered clients

authorize-server

Enables/disables the server's authorization on this client. If enabled, the server will send deny messages to the client that is trying to renew the lease, which the server knows the client shouldn't have

option82

Enable/disable the NAK of option 82 clients for this subnet

15.2.8.2.2. Client Configuration Options: Hosts

To set DHCP client configuration options at the host level, enter edit mode and navigate to */services/dhcpserver/host{host id}/options*.

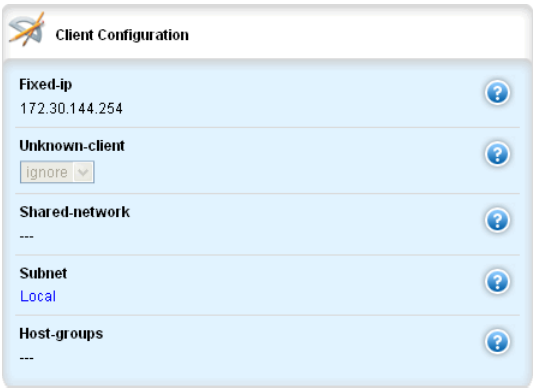

Figure 15.13. Client Configuration form for Hosts

fixed-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address that the server assigns to the matching client

unknown-client

Synopsis: *string - one of the following keywords { ignore, deny, allow }*

The action to take for previously unregistered clients

shared-network

Synopsis: *A string*

Shared-network that this host belongs to

subnet

Synopsis: *A string*

Subnet that this host belongs to

host-groups

Synopsis: *A string*

Host groups that this host belongs to

15.2.8.2.3. Client Configuration Options: Host-groups

To set DHCP client configuration options at the host-group level, enter edit mode and navigate to */services/dhcpserver/host-group{host-group id}/options*.

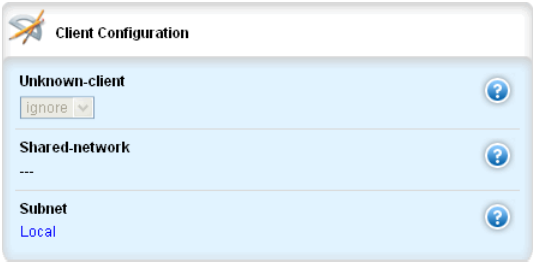

Figure 15.14. Client Configuration form for Host-groups

unknown-client

Synopsis: *string - one of the following keywords { ignore, deny, allow }*

Default: *allow*

The action to take for previously unregistered clients

shared-network

Synopsis: *A string*

Shared-network that this host group belongs to

subnet

Synopsis: *A string*

The subnet that this host group belongs to

15.2.8.3. Client Configuration Options at the DHCP Client Level

The following DHCP client, NIS, and NetBios options are set at the client level under all DHCP levels. To set the client configuration options, enter edit mode and navigate to:

- DHCP server options: */services/dhcpserver/options/client*
- subnet options: */services/dhcpserver/subnet{subnet id}/options/client*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options/client*
- host-group options: */services/dhcpserver/host-groups{host-group id}/options/client*
- host options: */services/dhcpserver/host{host id}/options/client*

Client Configuration:

Figure 15.15. Client Configuration form for DHCP Clients

hostname

Synopsis: *A string*

The unique name to refer to the host within a DHCP configuration

subnetmask

Synopsis: *IPv4 address in dotted-decimal notation*

Subnet mask

default-route

Synopsis: *IPv4 address in dotted-decimal notation*

The default route that the server offers to the client when it issues the lease to the client

broadcast

Synopsis: *IPv4 address in dotted-decimal notation*

The broadcast address that the server offers to the client when it issues the lease to the client

domain

Synopsis: *string*

The domain name that the server offers to the client when it issues the lease to the client

dns-server

Synopsis: *IPv4 address in dotted-decimal notation*

The domain name server that the server offers to client when it issues the lease to client

static-route

Synopsis: *IPv4 address in dotted-decimal notation*

The static route that the dhcpserver offers to the client when it issues the lease to the client

NIS Configuration:

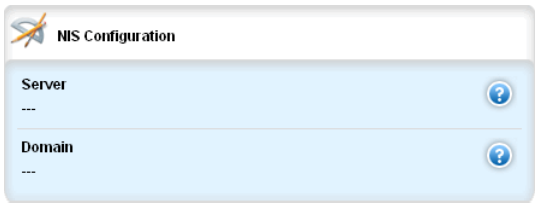

Figure 15.16. NIS Configuration form

server

Synopsis: *IPv4 address in dotted-decimal notation*

The NIS server address that the dhcpserver offers to the client when it issues the lease to the client

domain

Synopsis: *string*

The NIS domain name that the dhcpserver offers to the client when it issues the lease to the client

NetBios Configuration:

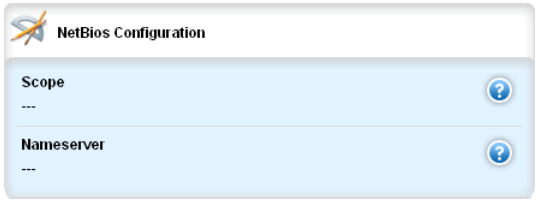

Figure 15.17. Netbios Configuration form

This is the NetBIOS nameserver that dhcpserver offers to client when it issues the lease to a client.

scope

Synopsis: *string*

Default: *netbios*

The NetBIOS scope that the dhcpserver offers to the client when it issues the lease to the client

nameserver

Synopsis: *string*

Default: *127.0.0.1*

The NetBIOS nameserver that the dhcpserver offers to the client when it issues the lease to the client

15.2.9. Custom DHCP Options

You can set custom DHCP options at the under clients at all DHCP levels. To set a custom DHCP option, you need to know the number of the option you want to set and the valid values for the option.

To set a custom DHCP option, enter edit mode and navigate to one of the following locations and click **<Add custom>**:

- DHCP server options: */services/dhcpserver/options/client/custom*
- subnet options: */services/dhcpserver/subnet{subnet id}/options/client/custom*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options/client/custom*
- host-group options: */services/dhcpserver/host-groups{host-group id}/options/client/custom*
- host options: */services/dhcpserver/host{host id}/options/client/custom*

On the Key settings form, set the number and value for the DHCP custom option.

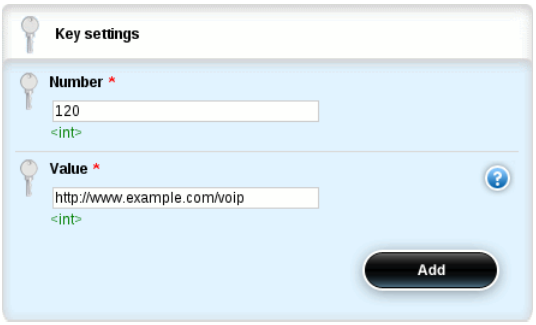

Figure 15.18. Setting a DHCP Custom Option

15.2.10. Hardware Configuration

The *Hardware Configuration* form is only available for DHCP hosts. To set the hardware options for a DHCP host, enter edit mode and navigate to */services/dhcpserver/host{host_01}/options*.

Figure 15.19. Hardware Configuration form

type

Synopsis: *string - one of the following keywords { ethernet, token-ring, fddi }*

Default: *ethernet*

The type of network hardware used by the client, associated with the host entry.

mac

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The physical network address of the client. Note that this corresponds to the hardware type; for example, MAC address for ethernet.

Part II. Network Interfaces and Ethernet Bridging

This part describes network interfaces and the configuration and monitoring of Ethernet bridging on a ROX™-based networking device:

Ethernet Ports	Chapter 16, *Ethernet Ports*
Ethernet Statistics	Chapter 17, *Ethernet Statistics*
IP Statistics	Chapter 18, *IP Statistics*
Virtualswitch Bridging	Chapter 19, *Virtual Switch Bridging*
Link Aggregation	Chapter 20, *Link Aggregation*
Modem	Chapter 21, *Modem*
Serial Ports	Chapter 22, *Serial Protocols*
WAN	Chapter 23, *WAN*
Port Security	Chapter 24, *Port Security*
Multicast Filtering	Chapter 25, *Multicast Filtering*
Classes of Service (CoS)	Chapter 26, *Classes Of Service*
MAC Address Tables	Chapter 27, *MAC Address Tables*
Spanning Tree Protocol	Chapter 28, *Spanning Tree*
Virtual LAN (VLAN)	Chapter 29, *Virtual LANs*
Network Discovery	Chapter 30, *Network Discovery*

16. Ethernet Ports

ROX™ Ethernet port control provides the following features:

- Configuring port physical parameters.
- Configuring link alarms/traps for the port.
- Configuring port rate limiting.
- Establishing port mirroring.
- Cable diagnostics.
- Viewing port status.
- Resetting one or more ports.
- Configuring Link-Fault-Indication (LFI).

16.1. Controller Protection Through Link-Fault-Indication (LFI)

Modern industrial controllers often feature backup Ethernet ports used in the event of a link failure. When these interfaces are supported by media (such as fiber) that employ separate transmit and receive paths, the interface can be vulnerable to failures that occur in only one of the two paths.

For example, refer to Figure 16.1, “Controller Protection Through LFI”. While the link between Switch A and the Controller functions normally, the Controller holds the backup link down. Switch B learns that to reach the Controller, it must forward frames towards Switch A.

If the transmission path from the Controller to Switch A fails, Switch A will still generate link signals to the Controller. The Controller will still detect the link to Switch A and will not fail over to the backup port.

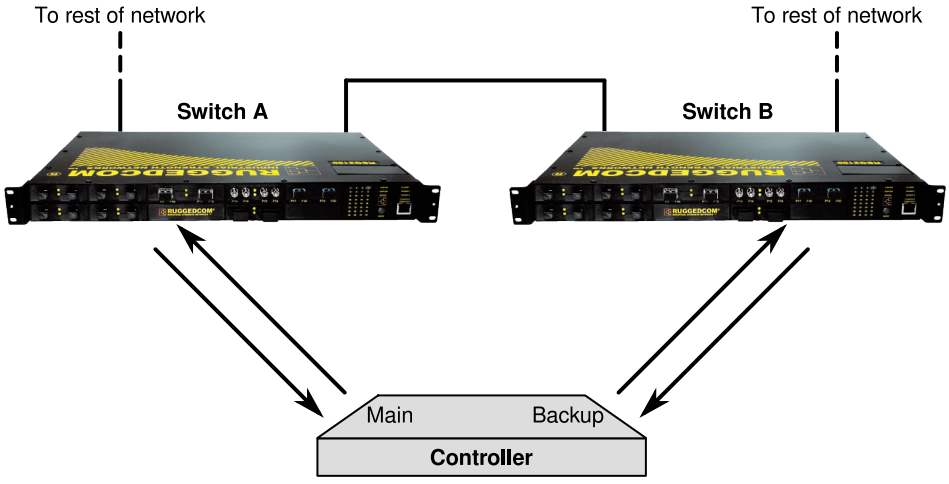

Figure 16.1. Controller Protection Through LFI

To overcome this problem, a way is needed to notify the link partner when receipt of the partner’s link integrity signal stops. Such a way exists natively in some link media but not in others:

- *Auto-Negotiating links (100Base-TX,1000Base-T,1000Base-X)* – auto-negotiation is a built-in feature; the Remote Fault Indication flag is set in the transmitted auto-negotiation signal.
- *100Base-FX links* – Far-End-Fault-Indication (FEFI) is a standard feature defined by the IEEE 802.3 standard for this link type. FEFI includes:

- Transmitting FEFI – transmits a modified link integrity signal when a link failure is detected (that is, when no link signal is received from the link partner).
- Detecting FEFI – indicates link loss when a FEFI signal is received from the link partner.

- *10Base-FL links* – no standard support.

10Base-FL links have no native link partner notification mechanism. Also, FEFI support in 100Base-FX links is optional according to the IEEE 802.3 standard, which means that some link partners may not support it.

RuggedCom offers an advanced Link-Fault-Indication (LFI) feature for the links where no native link partner notification mechanism is available. With LFI enabled, the device bases generation of a link integrity signal upon its reception of a link signal. In Figure 16.1, “Controller Protection Through LFI”, if Switch A fails to receive a link signal from the Controller, it will stop generating a link signal. The Controller will detect the link failure and switch to the backup port.

*If both link partners support LFI, LFI must **not** be enabled on both sides of the link. If LFI is enabled on both sides, the link will never be established because each side will permanently wait for its partner to transmit a link signal.*

16.2. Ethernet Port Configuration

The Ethernet Ports menu is accessible from the main menu under *interface*.

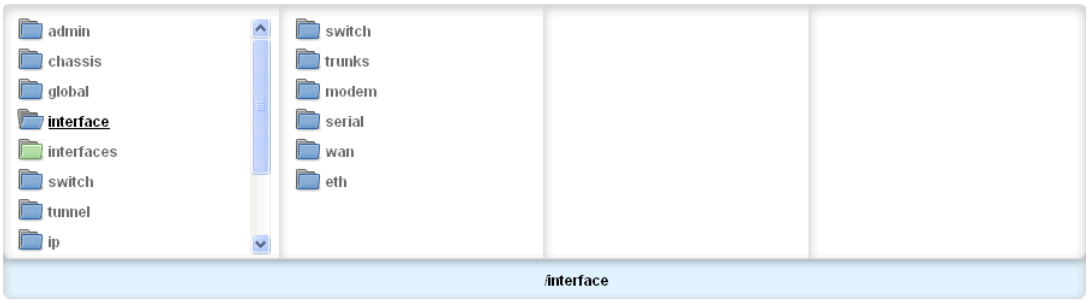

Figure 16.2. Ethernet Ports menu

16.2.1. Port Parameters

Switched Ethernet Ports

Next 16 (19)

Slot	Port	Enabled	Auto-Negotiation	Speed	Duplex	Flow Control	LFI
lm1	1	enabled	on	auto	auto	disabled	---
lm1	2	enabled	on	auto	auto	disabled	---
lm1	3	enabled	on	auto	auto	disabled	---
lm1	4	enabled	on	auto	auto	disabled	---
lm1	5	enabled	on	auto	auto	disabled	---
lm1	6	enabled	on	auto	auto	disabled	---
lm1	7	enabled	on	auto	auto	disabled	---
lm1	8	enabled	on	auto	auto	disabled	---
lm1	9	enabled	on	auto	auto	disabled	---
lm1	10	enabled	on	auto	auto	disabled	---
lm1	11	enabled	on	auto	auto	disabled	---
lm1	12	enabled	on	auto	auto	disabled	---
lm1	13	enabled	on	auto	auto	disabled	---
lm1	14	enabled	on	auto	auto	disabled	---
lm1	15	enabled	on	auto	auto	disabled	---
lm1	16	enabled	on	auto	auto	disabled	---

Figure 16.3. Switched Ethernet Ports table

The Switched Ethernet Ports table shows the Ethernet interfaces.

To display the Switched Ethernet Ports table, navigate to *interface/switch*.

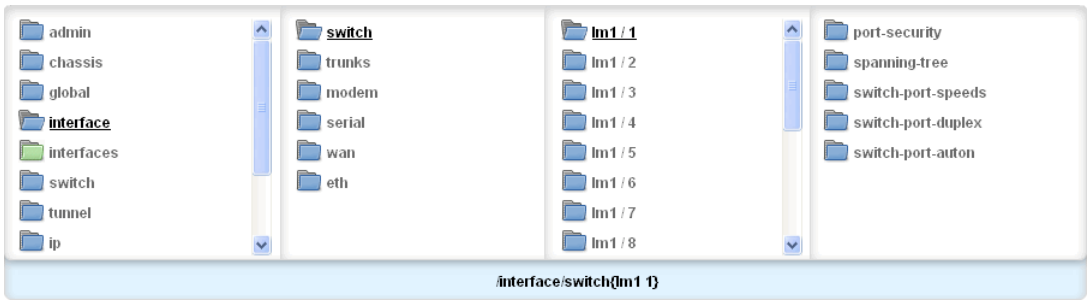

Figure 16.4. Switched Ethernet Ports submenu

The Switched Ethernet Ports Forms are accessible from submenus of the Ethernet Ports menu. To display the forms, navigate to *interface/switch/{line module}*.

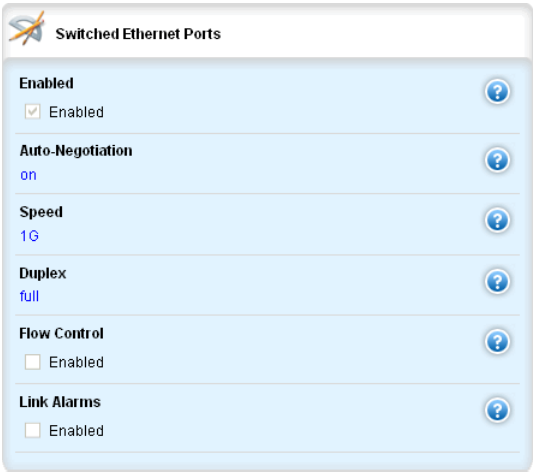

Figure 16.5. Switched Ethernet Ports form

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Enabled

Synopsis: *boolean*

Default: *true*

Provides the option to enable or disable this interface. When unchecked(i.e disabled), the interface will prevent all frames from being sent and received on that interface.

AutoN

Synopsis: *string - one of the following keywords { off, on }*

Enables or disables IEEE 802.3 auto-negotiation. Enabling auto-negotiation results in speed and duplex being negotiated upon link detection; both end devices must be auto-negotiation compliant for the best possible results.

Speed

Synopsis: *string - one of the following keywords { 10M, 100M, 1G, 10G, auto }*

Speed (in megabits-per-second or gigabits-per-second). If auto-negotiation is enabled, this is the speed capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is explicitly forced to this speed mode. AUTO means advertise all supported speed modes.

Duplex

Synopsis: *string - one of the following keywords { full, half, auto }*

If auto-negotiation is enabled, this is the duplex capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is explicitly forced to this duplex mode. AUTO means advertise all supported duplex modes.

Link Alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

Switchport

Synopsis: *boolean*

Default: *true*

Sets the physical port into either switched mode or a dedicated routing mode.

Flow Control

Flow control is useful for preventing frame loss during times of severe network traffic

LFI

Link Fault Indication (LFI) is specifically for FX interfaces.

Alias

Synopsis: *A string*

The SNMP alias name of the interface

> *If one end of the link is fixed to a specific speed and duplex type and the peer auto-negotiates, there is a strong possibility that the link will either fail to raise, or raise with the wrong settings on the auto-negotiating side. The auto-negotiating peer will fall back to half-duplex operation, even when the fixed side is full duplex. Full-duplex operation requires that both ends are configured as such or else severe frame loss will occur during heavy network traffic. At lower traffic volumes the link may display few if any errors. As the traffic volume rises, the fixed negotiation side will begin to experience dropped packets while the auto-negotiating side will experience excessive collisions. Ultimately, as traffic load approaches 100% the link will become entirely unusable. These problems can be avoided by always configuring ports to the appropriate fixed values.*

16.2.2. Port Rate Limiting

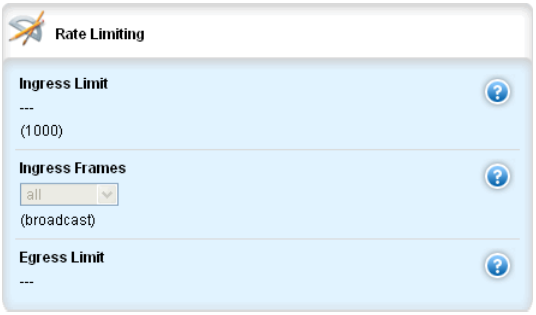

Figure 16.6. Rate Limiting form

To display the Rate Limiting forms, navigate to *interface/switch/{line module}*.

Ingress Limit

Synopsis: *string - the keyword { disabled }*

Synopsis: *integer*

Default: *1000*

The data rate in kbps at which received frames (of the type described by the ingress frames parameter) will start to be discarded by the switch. The valid range is 62 to 256000 kbps. The default value is 1000 kbps. If not set(cleared), this feature is disabled.

Ingress Frames

Synopsis: *string - one of the following keywords { all, mcast-flood-ucast, multicast, broadcast }*

Default: *broadcast*

This parameter specifies the types of frames to rate-limit on this port. It applies only to received frames:

- BROADCAST : only broadcast frames will be limited.
- MULTICAST : all multicast frames (including broadcast) will be limited.
- MCAST-FLOOD-UCAST : all multicast frames (including broadcast) will be limited. Unicast will not be limited.
- ALL : all frames (both multicast and unicast) will be limited.

Egress Limit

Synopsis: *string - the keyword { disabled }*

Synopsis: *integer*

Default: *disabled*

The maximum data rate in kbps at which the switch will transmit (multicast, broadcast and unicast) frames on this port. The switch will discard frames in order to meet this rate if required. The valid range is 62 to 256000 Kbps. If not set, this feature is disabled.

16.2.3. Port Mirroring

Port mirroring is a troubleshooting tool that copies, or mirrors, all traffic received or transmitted on a designated port to another mirror port. If a protocol analyzer were attached to the target port, the traffic stream of valid frames on any source port is made available for analysis.

Select a target port that has a higher speed than the source port. Mirroring a 100 Mbps port onto a 10 Mbps port may result in an improperly mirrored stream.

Frames will be dropped if the full-duplex rate of frames on the source port exceeds the transmission speed of the target port. Since both transmitted and received frames on the source port are mirrored to the target port, frames will be discarded if the sum traffic exceeds the target port's transmission rate. This problem reaches its extreme in the case where traffic on a 100 Mbps full-duplex port is mirrored onto a 10 Mbps half-duplex port.

Invalid frames received on the source port will not be mirrored. These include CRC errors, oversize and undersize packets, fragments, jabbers, collisions, late collisions and dropped events).

16.2.3.1. Port Mirroring Limitations

- Traffic will be mirrored onto the target port only if the target port is a member of the same VLANs as the source port.
- The target port may sometimes incorrectly show the VLAN tagged/untagged format of the mirrored frames.
- Network management frames (such as RSTP, GVRP etc.) may not be mirrored.
- Switch management frames generated by the switch (such as Telnet, HTTP, SNMP etc.) may not be mirrored.

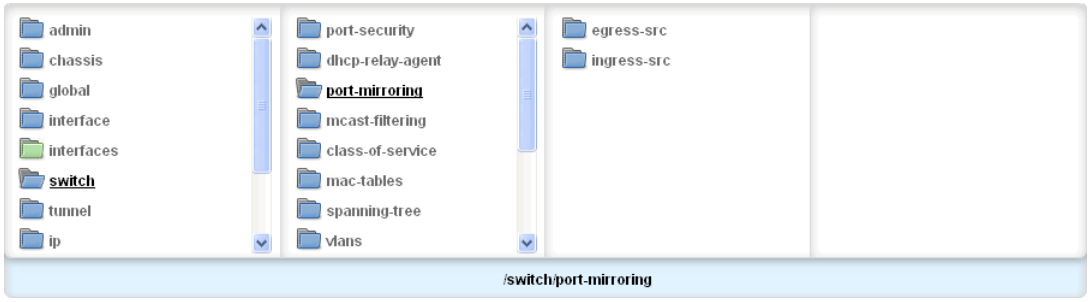

Figure 16.7. Port-Mirroring menu

To display the Port-Mirroring menu and Port Mirror form, navigate to *switch/port-mirroring*.

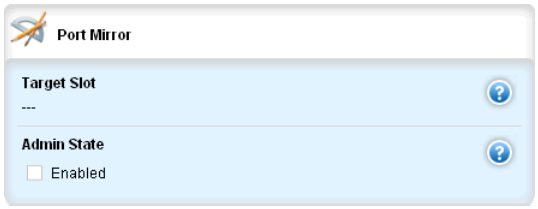

Figure 16.8. Port Mirror form

Target Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The slot where a monitoring device should be connected.

Target Port

Synopsis: *integer*

The port where a monitoring device should be connected.

Figure 16.9. Ingress Source Ports table

To display the Ingress Source Ports table, navigate to *switch/port-mirroring/ingress-src*.

Ingress Source Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Ingress Source Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

To display the Egress Source Ports table, navigate to *switch/port-mirroring/egress-src*.

Figure 16.10. Egress Source Ports table

Egress Source Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Egress Source Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

16.2.4. Diagnostics

To display the Diagnostics menu and Cable Diagnostics Results form, navigate to *interfaces/switch/{line module}/diagnostics*.

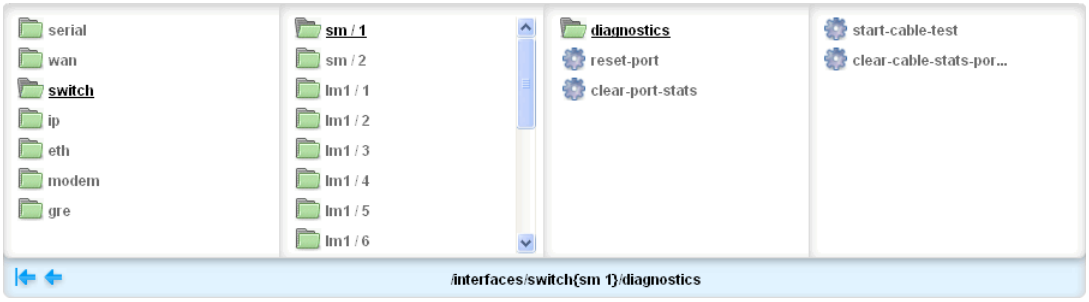

Figure 16.11. Diagnostics menu

ROX™ is able to perform cable diagnostics per Ethernet port and to view the results.

> *When cable diagnostics are performed on a port, any established network link on the port will be dropped and normal network traffic will not be able to pass through either the Port Under Test or the Partner Port. Please be aware of the potential network interruption that could be triggered by running cable diagnostics. After the cable diagnostics finish, the original network port settings for both the Port Under Test and the Partner Port are restored along with any established link.*

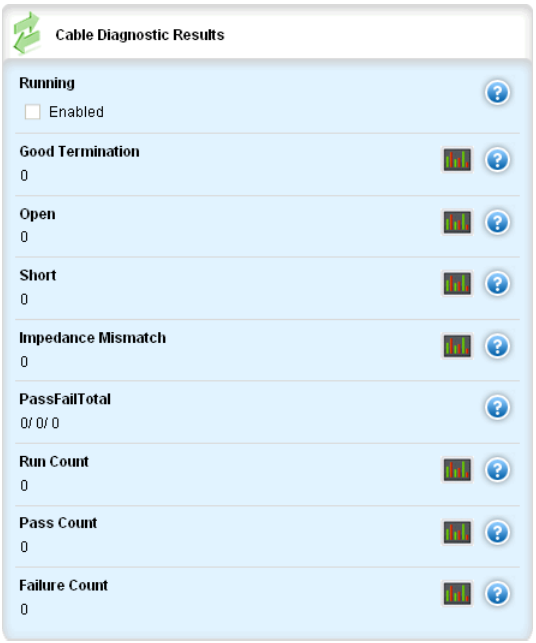

Figure 16.12. Cable Diagnostics Results form

Figure 16.12, “Cable Diagnostics Results form” displays the current value of diagnostic parameters for the corresponding Ethernet port. This form can be used to set certain cable diagnostic parameters for the port, as indicated below:

Running

Synopsis: *boolean*

Whether or not a cable test is currently running on this port

Good Termination

Synopsis: *unsigned short integer*

The number of times GOOD TERMINATION (no fault) is detected on the cable pairs of the selected port.

Open

Synopsis: *unsigned short integer*

The number of times OPEN is detected on the cable pairs of the selected port.

Short

Synopsis: *unsigned short integer*

The number of times SHORT is detected on the cable pairs of the selected port.

Impedance Mismatch

Synopsis: *unsigned short integer*

The number of times IMPEDANCE MISMATCH is detected on the cable pairs of the selected port.

PassFailTotal

Synopsis: *A string*

This field summarizes the results of the cable diagnostics performed so far.

- Pass : the number of times cable diagnostics were successfully completed on the selected port.
- Fail : the number of times cable diagnostics failed to complete on the selected port.
- Total : the total number of times cable diagnostics have been attempted on the selected port.

Run Count

Synopsis: *unsigned short integer*

Run Count : The total number of iterations

Pass Count

Synopsis: *unsigned short integer*

Pass Count

Failure Count

Synopsis: *unsigned short integer*

Failure Count

16.2.4.1. Running Cable Diagnostics

To start cable diagnostics on a port:

1. Connect a Category 5 or better quality cable to the port under test (PUT).
2. Connect the other end of the cable to a similar network port. For example, connect 100BASE-T port to a 100BASE-T port, 1000BASE-T port to a 1000BASE-T port.
3. Configure the PUT’s “Runs” count.
4. Configure the PUT’s cable diagnostics State to “Started”.

To stop cable diagnostics on a port:

1. Configure the PUT's cable diagnostics state to "Stopped". Diagnostics may be stopped at any point. If a stop is issued in the middle of a diagnostics run, it will nevertheless run to completion and the results will be updated.

> *Both the port under test (PUT) or partner port (PT) can be configured to be either in Enabled mode with auto-negotiation or in Disabled mode. Other modes may interfere with the cable diagnostics procedure and are not recommended.*

16.2.4.2. Interpreting Cable Diagnostics Results

Four different conditions are reported for the state of a cable under examination:

- Good – No fault is detected on the tested cable.
- Open – Opened cable pair(s) is/are detected on the tested cable.
- Short – Short cable pair(s) is/are detected on the tested cable.
- Imped – Impedance Mismatch is detected on the tested cable.

The corresponding counts for each of these status conditions indicates the number of occurrences of each type of fault. For a typical "no fault" Category 5 cable plugged into a 100BASE-T port, 'Good' will be incremented by two after every run of cable diagnostics, once for each cable pair used by a 100BASE-T port. Note that for a 1000BASE-T port, four cable pairs will be tested and so 'Good' will be incremented by four after every successful run.

For a fault condition, an estimated distance to the fault will be calculated and recorded in the system log. For detailed information about which cable pair has been detected to have experienced which type of fault and the corresponding distance to the fault, please refer to the system log file.

> *The "Runs" parameter cannot be changed while cable diagnostics are running on a port. In order to change the value, stop the diagnostic run on the port, change the "Runs" parameter, and restart diagnostics.*
>
> *On ports that do not support cable diagnostics, "N/A" will be shown as the cable diagnostics state and any settings made to the "Runs" and "Calibration" fields will be discarded.*

16.2.4.2.1. Cable Diagnostics Testing

To test cable diagnostics, navigate to *interfaces/switch/{line module}/diagnostics/start-cable-test* and follow the procedure outlined below.

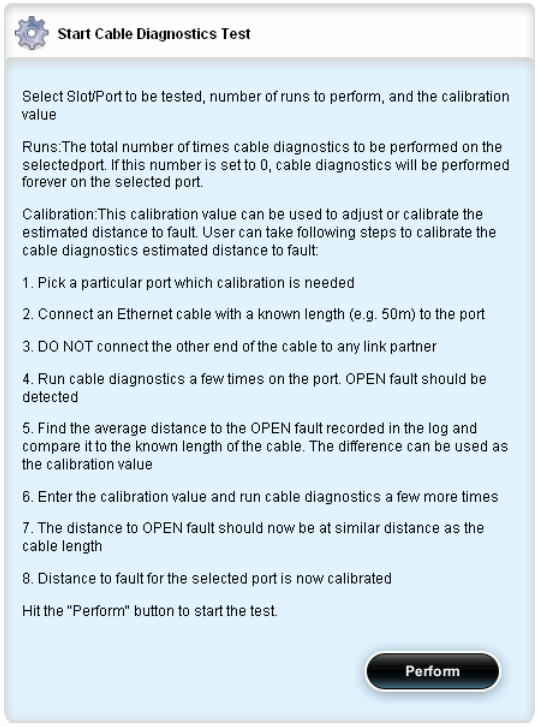
Start Cable Diagnostics Test
Select Slot/Port to be tested, number of runs to perform, and the calibration value
Runs:The total number of times cable diagnostics to be performed on the selectedport. If this number is set to 0, cable diagnostics will be performed forever on the selected port.
Calibration:This calibration value can be used to adjust or calibrate the estimated distance to fault. User can take following steps to calibrate the cable diagnostics estimated distance to fault:
1. Pick a particular port which calibration is needed
2. Connect an Ethernet cable with a known length (e.g. 50m) to the port
3. DO NOT connect the other end of the cable to any link partner
4. Run cable diagnostics a few times on the port. OPEN fault should be detected
5. Find the average distance to the OPEN fault recorded in the log and compare it to the known length of the cable. The difference can be used as the calibration value
6. Enter the calibration value and run cable diagnostics a few more times
7. The distance to OPEN fault should now be at similar distance as the cable length
8. Distance to fault for the selected port is now calibrated
Hit the "Perform" button to start the test.
Perform

Figure 16.13. Start Cable Diagnostics Test form

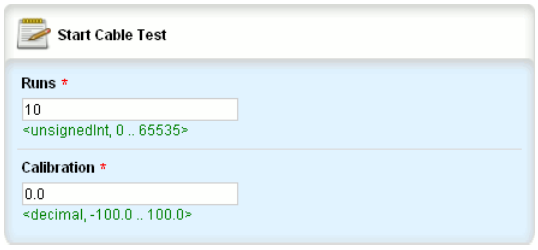

Figure 16.14. Start Cable Test form

To clear cable diagnostics, navigate to *interfaces/switch/{line module}/diagnostics/clear-cable-stats-port*. On the Clear Port Cable Diagnostic Test Results form, click **Perform**.

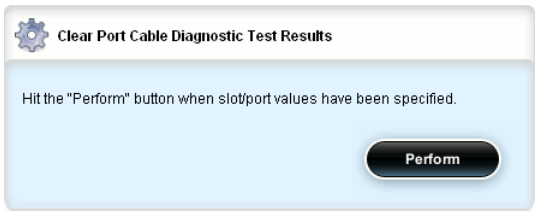

Figure 16.15. Clear Port Cable Diagnostic Test Results form

To clear all test results, rather than results from a single port, navigate to *switch/clear-cable-stats-all*.

Figure 16.16. Clear All Diagnostics (Switch) menu

To clear all cable diagnostic results, click the *Perform* button on the Clear All Cable Diagnostic Test Results form.

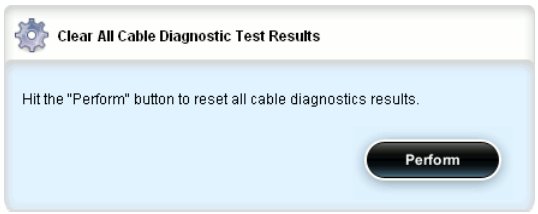

Figure 16.17. Clear All Cable Diagnostic Test Results form

16.2.4.2.2. Clearing Ethernet Alarms

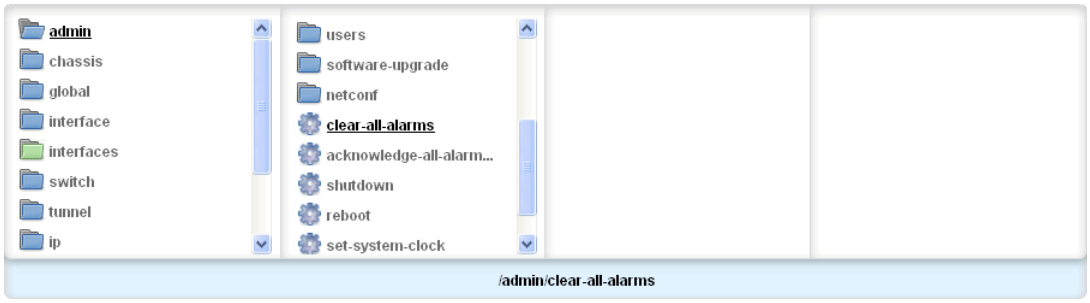

Figure 16.18. Clear All Alarms menu

Alarms can be cleared by hitting the *Perform* button. This command can be accessed from the Clear All Alarms menu action on the *admin/clear-all-alarms* menu.

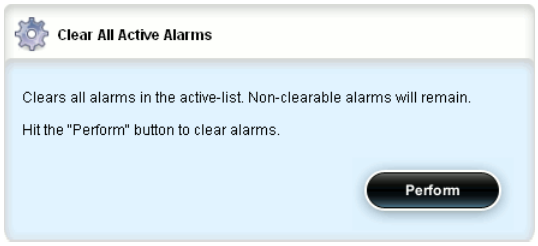
Clear All Active Alarms
Clears all alarms in the active-list. Non-clearable alarms will remain.
Hit the "Perform" button to clear alarms.
Perform

Figure 16.19. Clear All Active Alarms Trigger Action

16.2.4.3. Calibrating Estimated Distance To Fault

Follow these steps to set the “Calib” parameter (the estimated distance to fault):

1. Pick a particular port for which calibration is needed.
2. Connect an Ethernet cable with a known length (e.g. 50m) to the port.

3. Do not connect the other end of the cable to any link partner.
4. Run cable diagnostics a few times on the port. OPEN fault should be detected.
5. Find the average distance to the OPEN fault recorded in the log and compare it to the known length of the cable. The difference can be used as the calibration value.
6. Enter the calibration value and run cable diagnostics a few more times.
7. The distance to the OPEN fault should now be at a similar distance to the actual cable length.
8. The distance to the fault for the selected port is now calibrated.

16.2.5. Link Detection Options

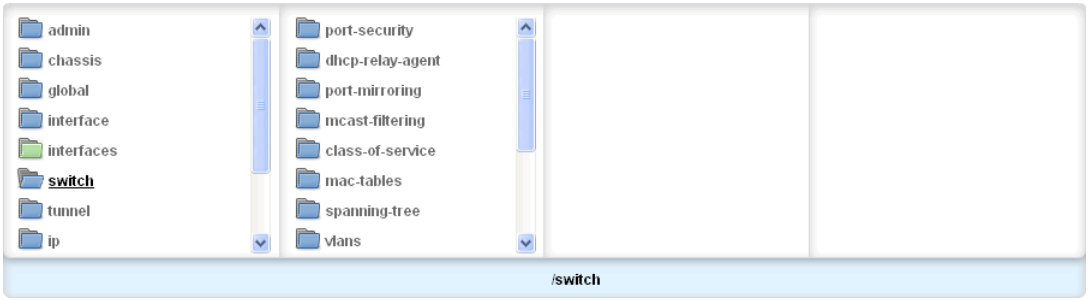

Figure 16.20. Switch (Link Detection) menu

The Switch (Link Detection) menu is accessible from the main menu under *switch*.

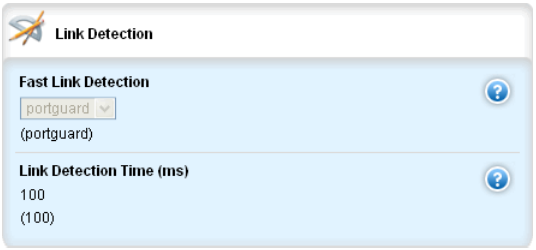

Figure 16.21. Link Detection form

The Link Detection form appears on the same screen as the *switch (Link Detection)* menu.

Fast Link Detection

Synopsis: *string - one of the following keywords { portguard, off, on }*

Default: *portguard*

Provides protection against faulty end devices generating an improper link integrity signal. When a faulty end device or a mis-matching fiber port is connected to the unit, a large number of continuous link state changes could be reported in a short period of time. This large number of bogus link state changes could render the system unresponsive as most, if not all, of the system resources are used to process the link state changes. This could in turn cause a serious network problem as the unit's RSTP process may not be able to run, thus allowing a network loop to form.

Three different settings are available for this parameter:

- ON_withPortGuard: This is the recommended setting. With this setting, an extended period (~2 minutes) of excessive link state changes reported by a port will prompt Port Guard feature to disable FAST LINK DETECTION on that port and raise an alarm. By disabling FAST LINK DETECTION on the problematic port, excessive link state changes can no longer consume a substantial amount of system resources. However if FAST LINK DETECTION is disabled, the port will need a longer time to detect a link failure. This may result in a longer network recovery

time of up to 2 seconds. Once Port Guard disables FAST LINK DETECTION on a particular port, the user can re-enable FAST LINK DETECTION on the port by clearing the alarm.

- ON: In certain special cases, where a prolonged excessive link state changes constitute a legitimate link operation, using this setting can prevent Port Guard from disabling FAST LINK DETECTION on the port in question. If excessive link state changes persist for more than 2 minutes, an alarm will be generated to warn the user about the observed bouncing link. If the condition of the excessive link state changes is resolved later on, the alarm will be cleared automatically. Since this option does not disable FAST LINK DETECTION, a persistent bouncing link could continue affect the system in terms of response time. This setting should be used with caution.
- OFF: Turning this parameter OFF will disable FAST LINK DETECTION completely. The switch will need a longer time to detect a link failure. This will result in a longer network recovery time of up to 2 seconds.

Link Detection Time (ms)

Synopsis: *integer*

Default: *100*

The time that the link has to continuously stay up before the 'link up' decision is made by the device. (The device performs de-bouncing of Ethernet link detection to avoid multiple responses to an occasional link bouncing event, e.g. when a cable is shaking while being plugged-in or unplugged).

When Fast Link Detection is enabled, the system prevents link state change processing from consuming all available CPU resources. If Port Guard is not used, however, it is possible for almost all available CPU time to be consumed by frequent link state changes, which could have a negative impact on overall system responsiveness.

16.3. Port Status

Figure 16.22. Interfaces menu

The *interfaces* menu is accessible from the main menu.

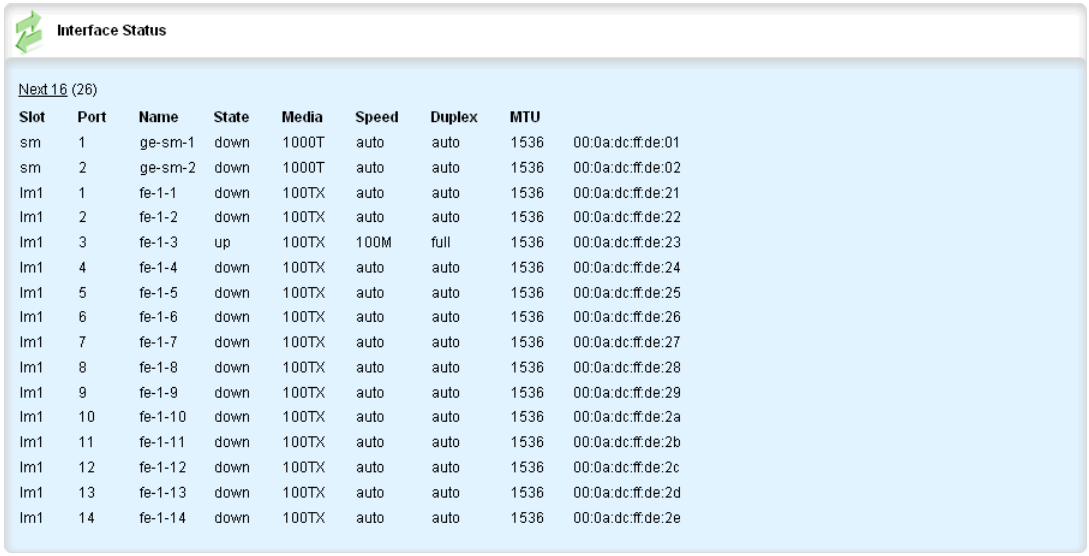
Interface Status
Next 16 (26)
Slot Port Name State Media Speed Duplex MTU
sm 1 ge-sm-1 down 1000T auto auto 1536 00:0a:dc:ff:de:01
sm 2 ge-sm-2 down 1000T auto auto 1536 00:0a:dc:ff:de:02
lm1 1 fe-1-1 down 100TX auto auto 1536 00:0a:dc:ff:de:21
lm1 2 fe-1-2 down 100TX auto auto 1536 00:0a:dc:ff:de:22
lm1 3 fe-1-3 up 100TX 100M full 1536 00:0a:dc:ff:de:23
lm1 4 fe-1-4 down 100TX auto auto 1536 00:0a:dc:ff:de:24
lm1 5 fe-1-5 down 100TX auto auto 1536 00:0a:dc:ff:de:25
lm1 6 fe-1-6 down 100TX auto auto 1536 00:0a:dc:ff:de:26
lm1 7 fe-1-7 down 100TX auto auto 1536 00:0a:dc:ff:de:27
lm1 8 fe-1-8 down 100TX auto auto 1536 00:0a:dc:ff:de:28
lm1 9 fe-1-9 down 100TX auto auto 1536 00:0a:dc:ff:de:29
lm1 10 fe-1-10 down 100TX auto auto 1536 00:0a:dc:ff:de:2a
lm1 11 fe-1-11 down 100TX auto auto 1536 00:0a:dc:ff:de:2b
lm1 12 fe-1-12 down 100TX auto auto 1536 00:0a:dc:ff:de:2c
lm1 13 fe-1-13 down 100TX auto auto 1536 00:0a:dc:ff:de:2d
lm1 14 fe-1-14 down 100TX auto auto 1536 00:0a:dc:ff:de:2e

Figure 16.23. Interface Status table

To display the Interface Status table, navigate to *interfaces/switch*.

Figure 16.24. Interface Status form

To display the Interface Status forms, navigate to *interfaces/switch/{line module}*.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the module.

Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

A descriptive name that may be used to identify the device connected on that port.

State

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

The port's link status.

Media

Synopsis: *A string*

The type of port media { 100TX, 10FL, 100FX, 1000X, 1000T, 802.11g, EoVDSL, 100TX }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

Speed

Synopsis: *string - one of the following keywords { 230.4K, 115.2K, 57.6K, 38.4K, 19.2K, 9.6K, 2.4K, 1.2K, 7.2M, 3.072M, 1.776M, 10G, 1G, 100M, 10M, 2.4M, 1.5M, auto }*

Speed (in Megabits-per-second or Gigabits-per-second)

Duplex

Synopsis: *string - one of the following keywords { full, half, auto }*

Duplex Setting: { Auto, Half, Full }

MTU

Synopsis: *integer*

The Maximum Transmission Unit of frame (in bytes) permitted on the interface.

MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The MAC Address of this specific port.

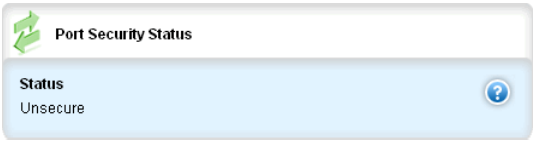

Figure 16.25. Port Security Status form

Security status describes the security status of the port.

Status

Synopsis: *A string*

The security status of the port.

16.4. Resetting Ports

This command performs a reset of the specified Ethernet ports. This action is useful for forcing re-negotiation of speed and duplex mode or in situations where the link partner has latched into an inappropriate state.

To reset ports, navigate to *interfaces/switch/{line module}/reset-port*. On the Reset Ethernet Port form, click *Perform*.

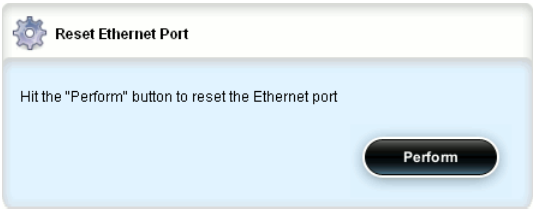

Figure 16.26. Reset Ethernet Port form

16.4.1. Resetting All Switched Ports

To reset all switched ports, navigate to *switch/reset-all-switched-ports*. On the Reset All Switched Ports form, click *Perform*.

Figure 16.27. Reset All Switched Ports menu

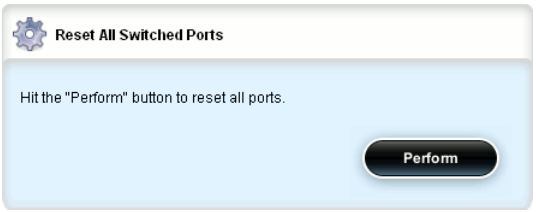

Figure 16.28. Reset All Switched Ports form

16.5. Troubleshooting

Problem One

One of my links seems to be fine at low traffic levels, but starts to fail as traffic rates increase.

One of my links pings OK but has problems with FTP/SQL/HTTP/…

A possible cause of intermittent operation is that of a 'duplex mismatch'. If one end of the link is fixed to full-duplex and the peer auto-negotiates, the auto-negotiating end falls back to half-duplex operation. At lower traffic volumes, the link may display few if any errors. As the traffic volume rises, the fixed negotiation side will begin to experience dropped packets while the auto-negotiating side will experience collisions. Ultimately, as traffic loads approach 100%, the link will become entirely unusable.

> *The ping command with flood options is a useful tool for testing commissioned links. The command "ping 192.168.0.1 500 2" can be used to issue 500 pings each separated by two milliseconds to the next switch. If the link used is of high quality, then no pings should be lost and the average round trip time should be small.*

Problem Two

I am trying to use the LFI protection feature but my links won't even come up.

Is it possible that the peer also has LFI enabled? If both sides of the link have LFI enabled, then both sides will withhold link signal generation from each other.

17. Ethernet Statistics

ROX™ provides the following features for gathering and reporting Ethernet statistics:

- Viewing basic Ethernet statistics.
- Viewing and clearing detailed Ethernet statistics.

The Ethernet Statistics menus are accessible from the main menu. The path to these menus is *interfaces/switch* and then clicking on any of the linked submenus from lm1/1 to lm1/14.

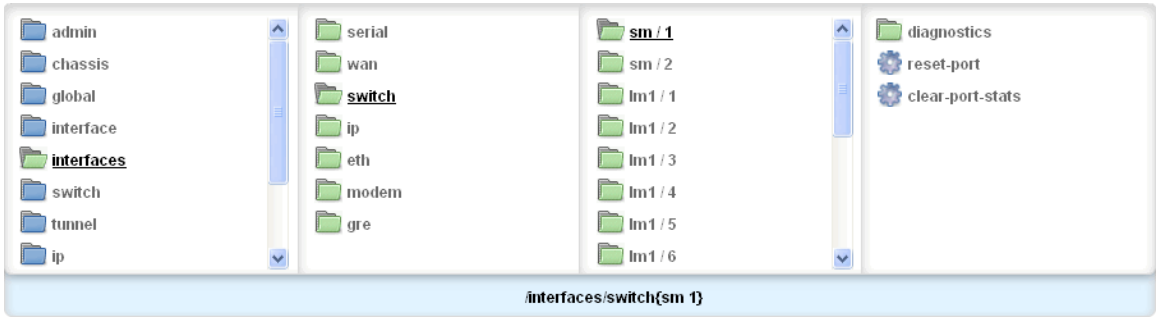

Figure 17.1. Ethernet Port Statistics Menu

17.1. Viewing Ethernet Statistics

This table provides basic Ethernet statistics information which is reset periodically, every few seconds. This traffic view is useful when the origin and destination of a traffic flow need to be determined.

17.2. Viewing Ethernet Port Statistics

Ethernet port statistics provide a detailed view of the traffic. This is useful when the exact source of error or traffic mix needs to be determined.

To display the ethernet port statistics forms, navigate to *interfaces/switch/{line module}*

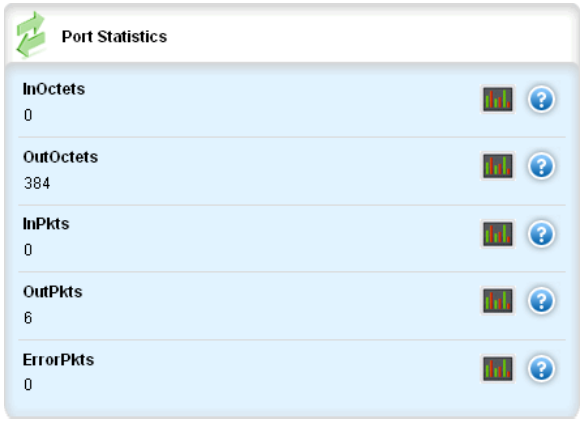

Figure 17.2. Port Statistics Form

InOctets

Synopsis: *unsigned integer*

The number of octets in received good packets. (Unicast+Multicast+Broadcast) and dropped packets.

OutOctets

Synopsis: *unsigned integer*

The number of octets in transmitted good packets.

InPkts

Synopsis: *unsigned integer*

The number of received good packets (Unicast+Multicast+Broadcast) and dropped packets.

OutPkts

Synopsis: *unsigned integer*

The number of transmitted good packets.

ErrorPkts

Synopsis: *unsigned integer*

The number of any type of erroneous packets.

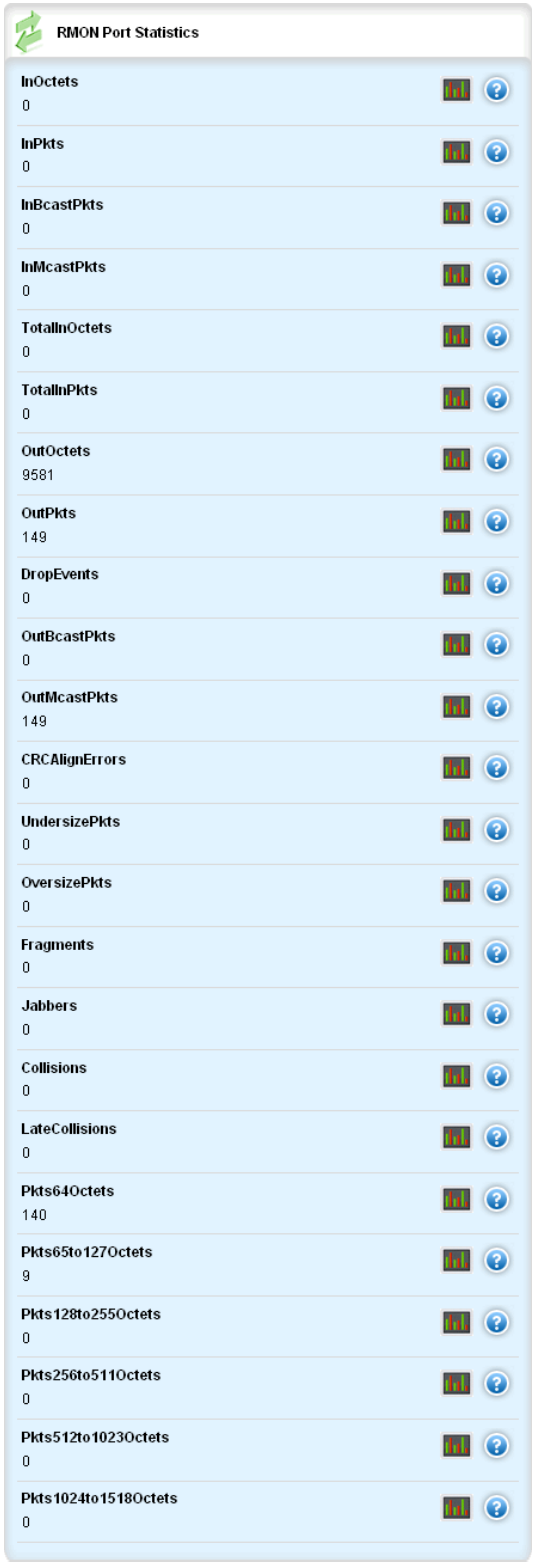

Figure 17.3. RMON Port Statistics Form

InOctets

Synopsis: *unsigned long integer*

The number of octets in received good packets (Unicast+Multicast+Broadcast) and dropped packets.

InPkts

Synopsis: *unsigned long integer*

The number of received good packets (Unicast+Multicast+Broadcast) and dropped packets.

InBcastPkts

Synopsis: *unsigned long integer*

The number of good broadcast packets received.

InMcastPkts

Synopsis: *unsigned long integer*

The number of good multicast packets received.

TotalInOctets

Synopsis: *unsigned long integer*

The total number of octets of all received packets. This includes data octets of rejected and local packets which are not forwarded to the switching core for transmission. It should reflect all the data octets received on the line.

TotalInPkts

Synopsis: *unsigned long integer*

The number of received packets. This includes rejected, dropped and local packets, as well as packets which are not forwarded to the switching core for transmission. It should reflect all packets received on the line.

OutOctets

Synopsis: *unsigned long integer*

The number of octets in transmitted good packets.

OutPkts

Synopsis: *unsigned long integer*

The number of transmitted good packets.

DropEvents

Synopsis: *unsigned integer*

The number of received packets that are dropped due to lack of receive buffers.

OutBcastPkts

Synopsis: *unsigned long integer*

The number of transmitted broadcast packets.

OutMcastPkts

Synopsis: *unsigned long integer*

The number of transmitted multicast packets. This does not include broadcast packets.

CRCAlignErrors

Synopsis: *unsigned integer*

The number of packets received which meet all the following conditions:

1. The packet data length is between 64 and 1536 octets inclusive.

2. The packet has invalid CRC.
3. A Collision Event has not been detected.
4. A Late Collision Event has not been detected.

UndersizePkts

Synopsis: *unsigned long integer*

The number of received packets which meet all the following conditions:

1. The packet data length is less than 64 octets.
2. A Collision Event has not been detected.
3. A Late Collision Event has not been detected.
4. The packet has valid CRC.

OversizePkts

Synopsis: *unsigned integer*

The number of packets received with data length greater than 1536 octets and valid CRC.

Fragments

Synopsis: *unsigned integer*

The number of packets received which meet all the following conditions:

1. The packet data length is less than 64 octets, or it is a packet without SFD and is less than 64 octets in length.
2. A Collision Event has not been detected.
3. A Late Collision Event has not been detected.
4. The packet has invalid CRC.

Jabbers

Synopsis: *unsigned integer*

The number of packets which meet all the following conditions:

1. The packet data length is greater that 1536 octets.
2. The packet has invalid CRC.

Collisions

Synopsis: *unsigned integer*

The number of received packets for which a Collision Event has been detected.

LateCollisions

Synopsis: *unsigned integer*

The number of received packets for which a Late Collision Event has been detected.

Pkts64Octets

Synopsis: *unsigned integer*

The number of received and transmitted packets with a size of 64 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets.

Pkts65to127Octets

Synopsis: *unsigned integer*

The number of received and transmitted packets with a size of 65 to 127 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets

Pkts128to255Octets

Synopsis: *unsigned integer*

The number of received and transmitted packets with a size of 128 to 257 octets. This includes received and transmitted packets as well as dropped and local received packets. This does not include rejected received packets

Pkts256to511Octets

Synopsis: *unsigned integer*

The number of received and transmitted packets with size of 256 to 511 octets.

This includes received and transmitted packets as well as dropped and local

received packets. This does not include rejected received packets.

Pkts512to1023Octets

Synopsis: *unsigned integer*

The number of received and transmitted packets with size of 512 to 1023 octets.

This includes received and transmitted packets as well as dropped and local

received packets. This does not include rejected received packets

Pkts1024to1518Octets

Synopsis: *unsigned integer*

The number of received and transmitted packets with a size of 1024 to 1536

octets. This includes received and transmitted packets as well as dropped and

local received packets. This does not include rejected received packets.

17.3. Viewing Non-switched Ethernet Statistics

The Non-switched Ethernet Statistics menus are accessible from the main menu under *Interfaces*. Statistics can be found under the submenus that follow. For instance, click on *eth* and then click on a linked submenu for an interface (for example, fe-cm-1).

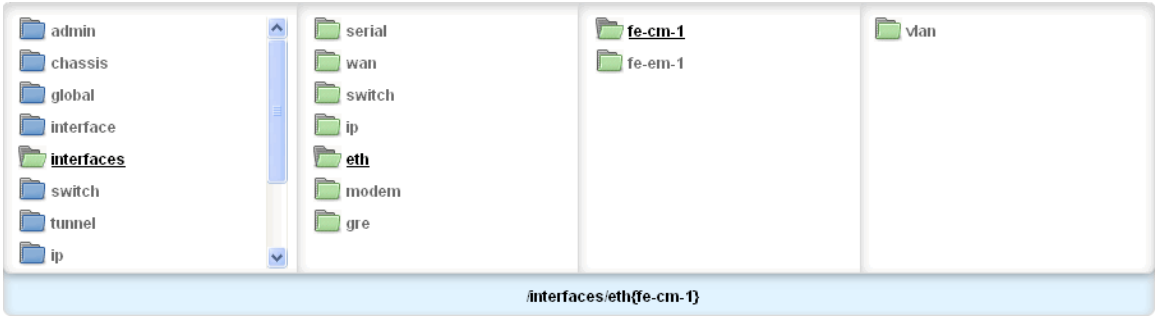

Figure 17.4. Statistics Menu

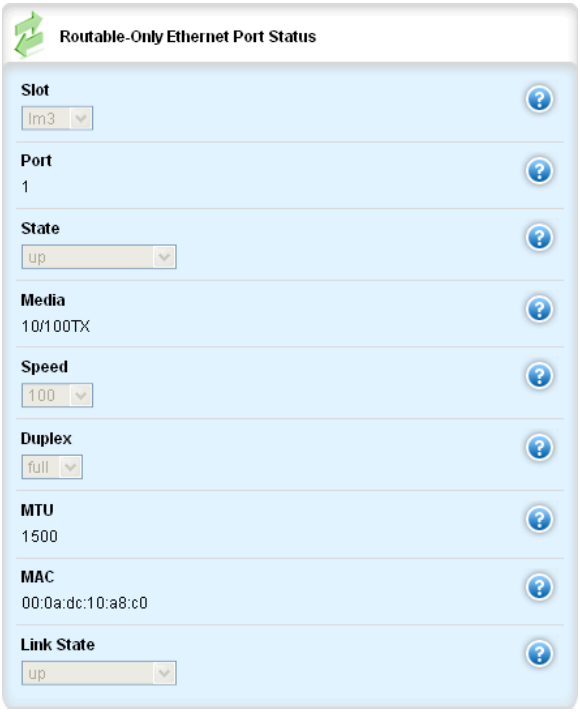

Figure 17.5. Routable-Only Ethernet Port Status Form

The Routable-Only Ethernet Port Status, Receive Statistics, and Transmit Statistics forms appear on the same screen as the *Statistics* menus. The Routable Ethernet Ports form displays the ethernet port configuration and status for a port. Ethernet statistics for the system's IP interfaces are available on the Receive Statistics and Transmit Statistics forms.

Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

Slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number on the module.

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

The port's link status.

media

Synopsis: *A string*

The type of port media { 100TX, 10FL, 100FX, 1000X, 1000T, 802.11g, EoVDSL, 100TX }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short

Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

Speed

Synopsis: *string - one of the following keywords { 1000, 100, 10 }*

Link speed (in Megabits-per-second or Gigabits-per-second)

Duplex

Synopsis: *string - one of the following keywords { full, half }*

Link duplex status.

MTU

Synopsis: *integer*

MTU (Maximum Transmission Unit) value on the port.

MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The MAC address of the port.

Link State

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

Link status.

Figure 17.6. Receive Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

Figure 17.7. Transmit Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

17.4. Clearing Switched Ethernet Port Statistics

To clear the switched ethernet port statistics, navigate to *interfaces/switch/{line module}/clear-port-stats*.

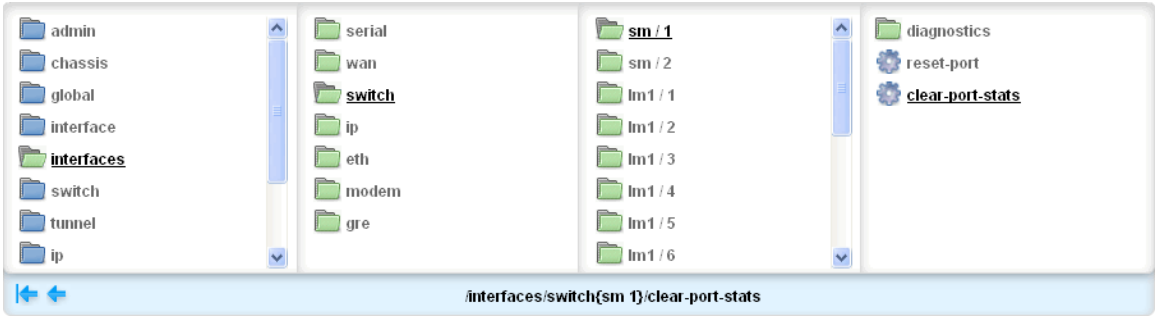

Figure 17.8. Interfaces Switch (Clearing Port Statistics) Menu

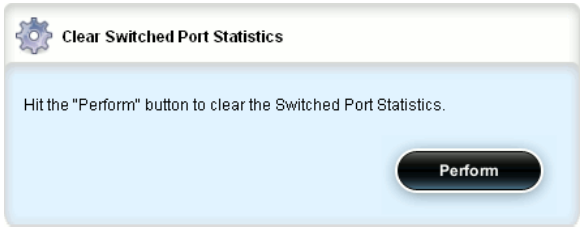

Figure 17.9. Clear Switched Port Statistics Form

This command clears Ethernet ports statistics for one switched port. Ports are cleared by clicking the *Perform* button on the Clear Switched Port Statistics form.

Figure 17.10. Clear All Statistics Menu

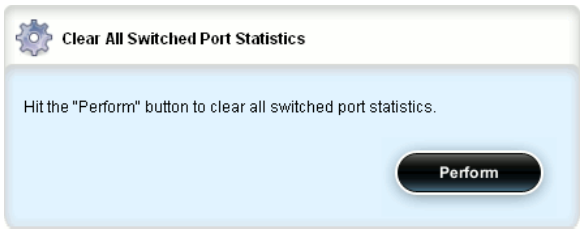

Figure 17.11. Clear All Switched Port Statistics Form

The *Clear All Switched Port Statistics* form clears all statistics for switched ethernet ports. To display the form, navigate to *switch/clear-all-switch-stats*.

18. IP Statistics

The forms and tables accessible from the Interfaces IP menu (below) show the status of what has been configured using the forms and tables from the Interface and IP menus.

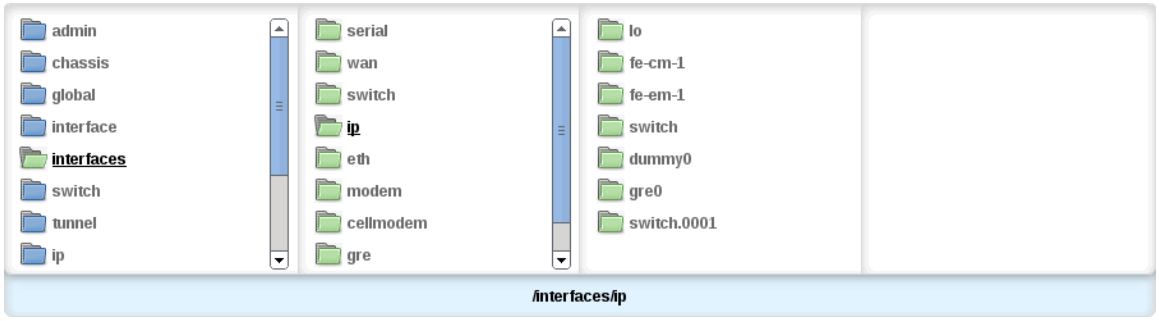

Figure 18.1. Interfaces IP Menu

The Interfaces IP menu is accessible from the main menu under *interfaces/ip*.

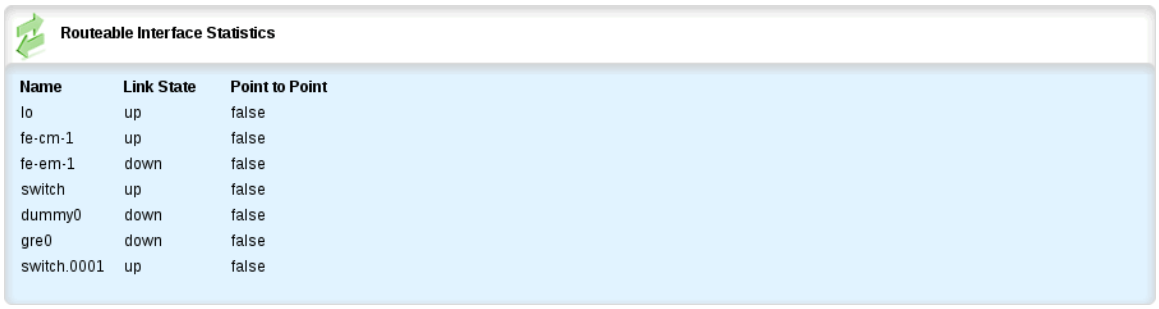

Figure 18.2. Routable Interface Statistics Table

This table appears on the same screen as the *Interfaces IP* menu.

The path to the Routable Interface Statistics form, Receive Statistics form and Transmit Statistics form is *interfaces/ip/{interface}*.

Figure 18.3. Routable Interface Statistics Form

Name

Synopsis: *A string*

The interface name

Link State

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

Up or Down

Point to Point

Synopsis: *boolean*

Is point to point link.

Figure 18.4. Receive Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

Figure 18.5. Transmit Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

19. Virtual Switch Bridging

19.1. Overview

A virtual switch bridges different network segments in way that is not dependent on a particular protocol.

Network traffic between segments is forwarded regardless of the IP and MAC addresses in a packet. In a virtual switch, forwarding is done in Layer 2 and allows all network traffic, including L2 Multicast (GOOSE, ISO), IP Multicast, Unicast, and Broadcast messages, to go through the virtual switch tunnel without any modifications. A virtual switch can be useful for GOOSE messaging when the sender and receiver need to communicate through a routable IP network. Because there is no IP encapsulation for the L2 traffic going through the virtual switch, network latency is minimized for the traffic between end devices.

The virtual switch appears on the device as a virtual Ethernet interface over a physical interface (T1/E1-Eth-HDLC, Ethernet port) between two routers. Physically, the two routers can be in different locations. There can be multiple virtual switch instances in a router. Each instance can include two or more interfaces, but an interface can only be a member of one virtual switch instance.

There can be multiple virtual switch interfaces over a "T1/E1 Eth-HDLC" interface in which the virtual switch interfaces are separated by creating a VLAN over the T1/E1 Eth-HDLC interface.

A virtual switch interface in a router can be a routable interface when an IP address is assigned either statically or via DHCP. The network address assigned to the virtual switch interface can be included in the dynamic routing protocol and the interface can carry a routing update. The IP address assigned to the virtual switch can be used as the default gateway for the end devices connected to the virtual switch interface. Network services, such as SSH, DHCP, NTP, VRRP, etc, can be configured to run on the virtual switch interface.

19.1.1. Helpful Hints

- Be careful when adding a VLAN interface (assigned to a switch port on a given line module) in the virtual switch. The VLAN tag on a tagged frame received on the VLAN Interface of a switch port may not be preserved when the traffic is egressed through a routable interface (T1/E1 Eth-HDLC, FE-CM-1) which is also part of the same virtual switch instance. However, a VLAN tag is preserved when tagged traffic is received on a routable interface. See Section 19.2, "Sample Use Case" for information on configuring a virtual switch that includes a switch port and a router port.
- Any IP address assigned to an interface becomes inactive and hidden when the interface is added to the virtual switch. The address on the interface is reactivated after removing the interface from the virtual switch.
- Be careful when adding interfaces to the virtual switch. Any network services running on the individual interfaces will need to be reconfigured after adding the interface to the virtual switch. For example, if a DHCP server running on FE-CM-1 is subsequently made a member of the VirtualSwitch VS1, the DHCP configuration must be changed to refer to VS1.
- In ROX™, the virtual switch is implemented in the software. Therefore, a CPU resource is needed to perform forwarding of broadcast, multicast and unicast traffic.
- If the router is running as a firewall, the routeback option must be enabled for the virtual switch interface in the "fwinterface" submenu under the Firewall menu.

19.2. Sample Use Case

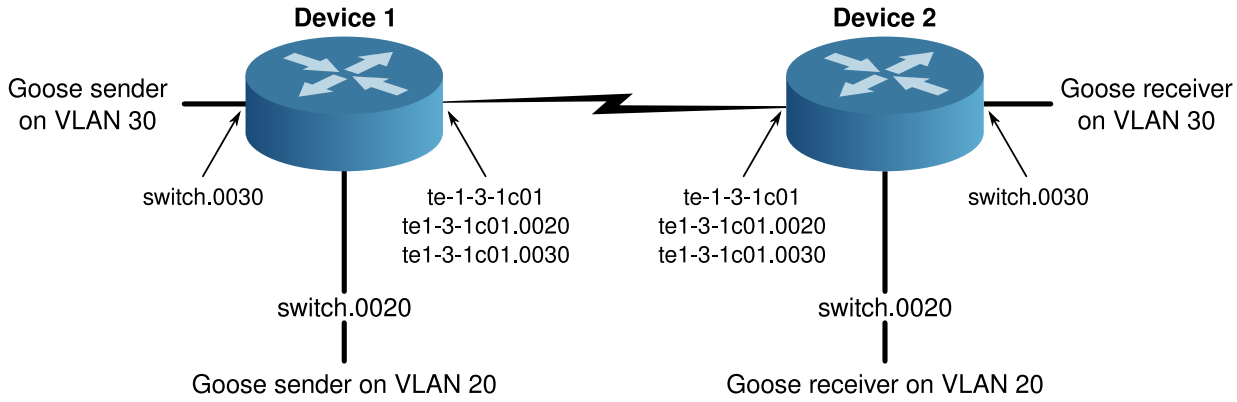

Figure 19.1. Virtual switch with multiple interfaces

To create the configuration shown in this example, follow these steps:

1. Configure the port connected to the senders and receivers as follows:
   - PVID 20, format as tagged.
   - PVID 30, format as tagged.
2. Configure hdlc-eth over T1/E1 on both devices with two VLANs: VLAN 20 and VLAN 30.
3. Configure two instances of VirtualSwitch by adding the following interfaces to the virtual switch on both devices:
   - VS1 on Device 1: switch.0020, te1-3-1c01.0020
   - VS2 on Device 1: switch.0030, te1-3-1c01.0030
4. Use the same configuration for Device 2.
5. Assign IP addresses to the virtual switch instances on both the devices:
   - VS1 on Device 1: 192.168.11.11/24
   - VS2 on Device 1: 192.168.22.22/24
   - VS1 on Device 2: 192.168.11.12/24
   - VS2 on Device 2: 192.168.22.23/24

When configuration is complete, tagged or untagged traffic received on VS1 of Device 1 should only be forwarded to VS1 on Device 2. Similarly, traffic received on VS2 of Device 1 should only be forwarded to VS2 on Device 2.

19.3. Virtual Switch Configuration and Status

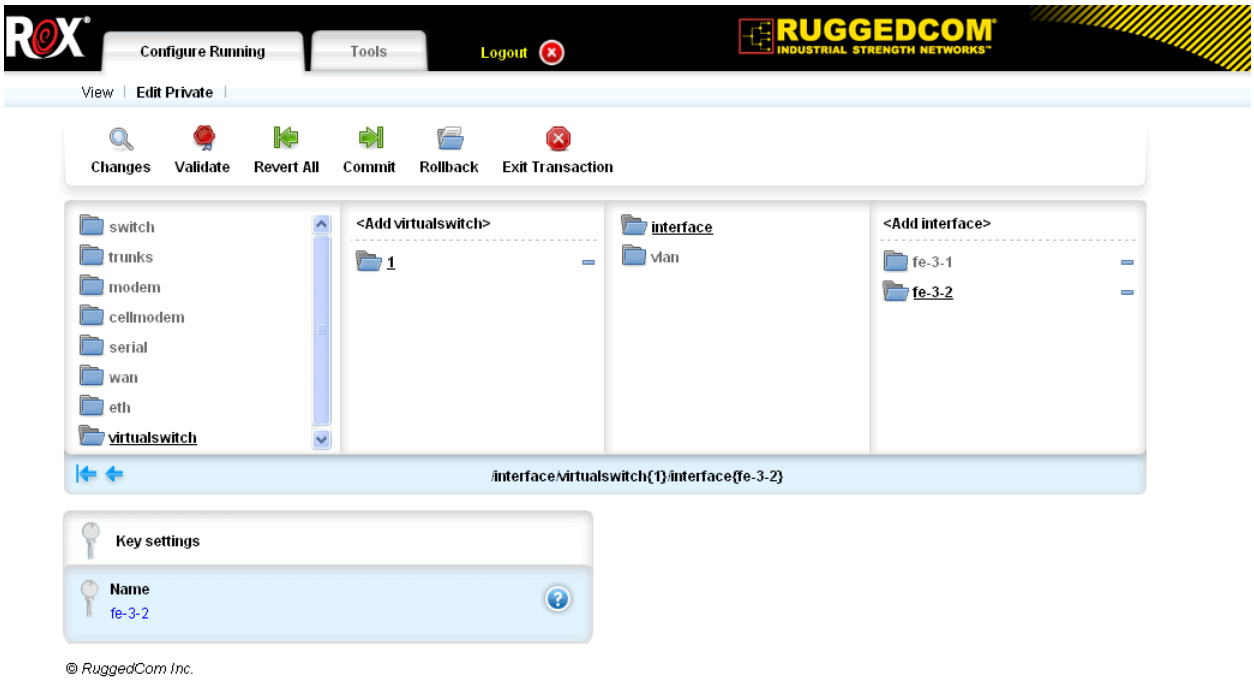

Figure 19.2. Adding a Virtual Switch

To add a virtual switch, enter Edit Private mode. Add a virtual switch and at least two interfaces. You can also add VLANs.

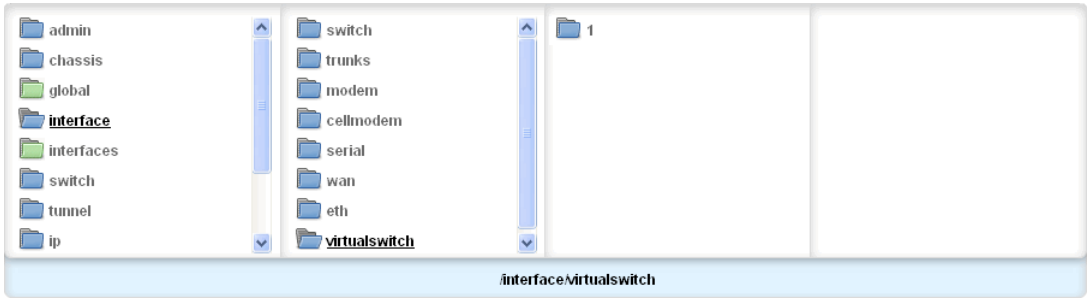

Figure 19.3. Interface Virtualswitch menu

The Interface Virtualswitch menu is located at *interface/virtualswitch*. If a virtual switch is configured, the Virtualswitch table appears on the same screen as this menu.

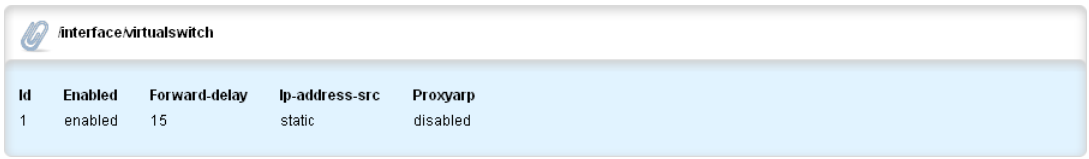

Figure 19.4. Virtualswitch table

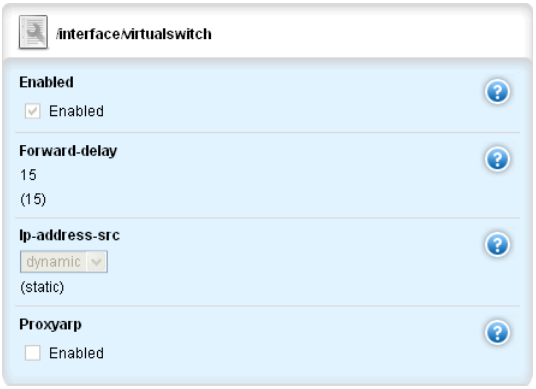

Figure 19.5. Virtualswitch form

To display this form, navigate to *interface/virtualswitch/{number}*.

Forward Delay

Synopsis: *unsigned byte*

Default: *15*

Delay (in seconds) of the listening and learning state before goes to forwarding state.

Alias

Synopsis: *A string*

The SNMP alias name of the interface

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP.

ProxyARP

Enables/Disables whether the port will respond to ARP requests for hosts other than itself

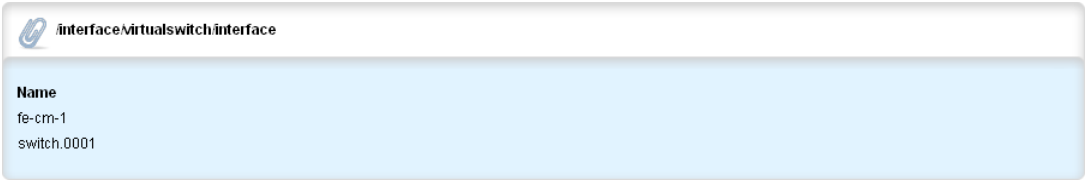

Figure 19.6. Interface table

To display this table, navigate to *interface/virtualswitch/{number}/interface*.

Interface Name

Synopsis: *A string*

Interface name.

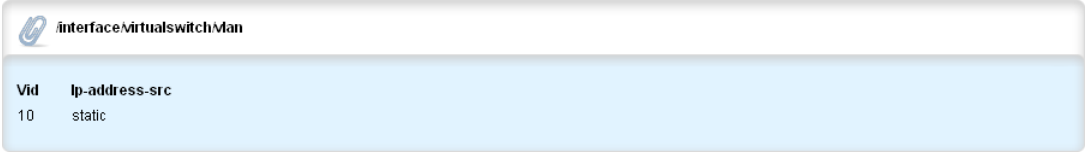

Figure 19.7. VLAN table

To display this table, navigate to *interface/virtualswitch/{number}/vlan*.

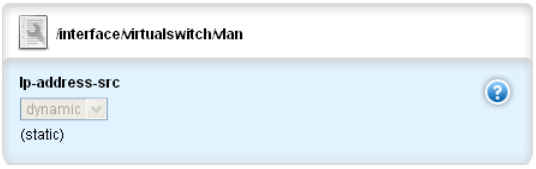

Figure 19.8. VLAN form

To display this form, navigate to *interface/virtualswitch/{number}/vlan/{number}*.

VLAN ID

Synopsis: *integer*

VLAN ID for this routable logical interface

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP.

If a virtual switch has been configured, some virtual switch data will be displayed under the Interfaces Virtualswitch menu.

Figure 19.9. Interfaces Virtualswitch menu

To display the Interfaces Virtualswitch menu, navigate to *interfaces/virtualswitch*. The Virtualswitch table appears on the same screen as this menu.

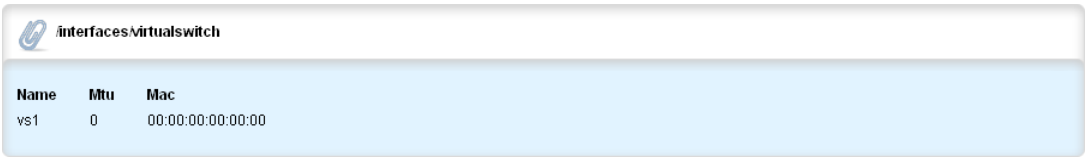

Figure 19.10. Virtualswitch table

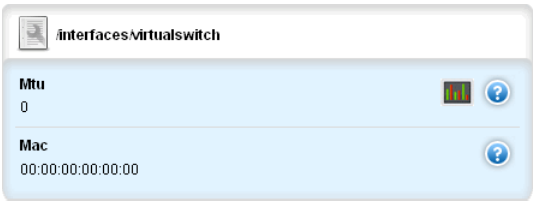

Figure 19.11. Virtualswitch form

To display the Virtualswitch, Receive, and Transmit forms, navigate to *interfaces/virtualswitch/{virtualswitch number}*.

Interface Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

MTU

Synopsis: *integer*

MTU (Maximum Transmission Unit) value on the port.

MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The MAC address of the port.

Figure 19.12. Receive form

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

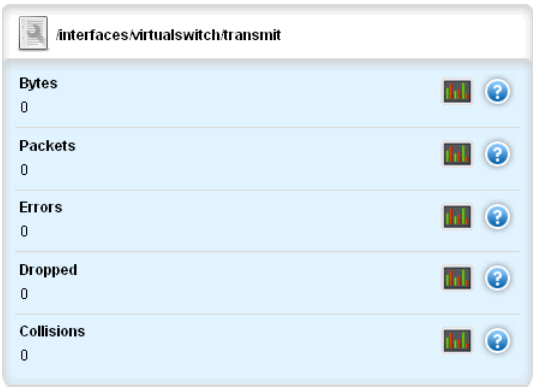

Figure 19.13. Transmit form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

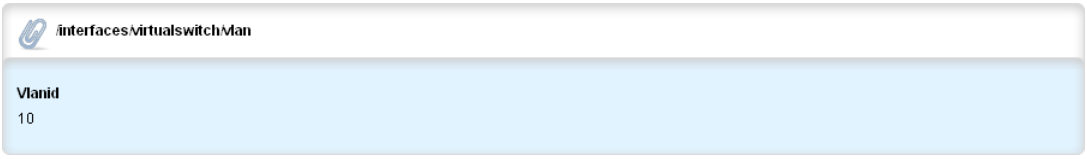

Figure 19.14. VLAN table

To display this table, navigate to *interfaces/virtualswitch/{virtualswitch number}/vlan*.

VLAN ID

Synopsis: *integer*

VLAN ID.

Figure 19.15. VLAN Receive form

To display the Receive and Transmit forms, navigate to *interfaces/virtualswitch/{virtualswitch number}/vlan/{number}*.

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped Into Abyss

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

Figure 19.16. VLAN Transmit form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

20. Link Aggregation

Link Aggregation aggregates or bundles several Ethernet ports into one logical link, called a *port trunk*, with higher bandwidth. Link Aggregation is also known as *port trunking* or *port bundling*.

ROX™ provides the following Link Aggregation features:

- Support for up to 15 port trunks.

> *The actual maximum number of port trunks depends on the number of ports in the switch (at least two ports are required to compose a port trunk).*

- Aggregation of up to 8 ports into a single port trunk.
- Highly randomized load balancing between the aggregated links, based on both the source and destination MAC addresses of the forwarded frames.

20.1. Link Aggregation Operation

Link Aggregation can be used for two purposes:

- To obtain increased and linearly incremental link bandwidth.
- To improve network reliability by creating link redundancy. If one of the aggregated links fails, the switch will balance the traffic between the remaining links.

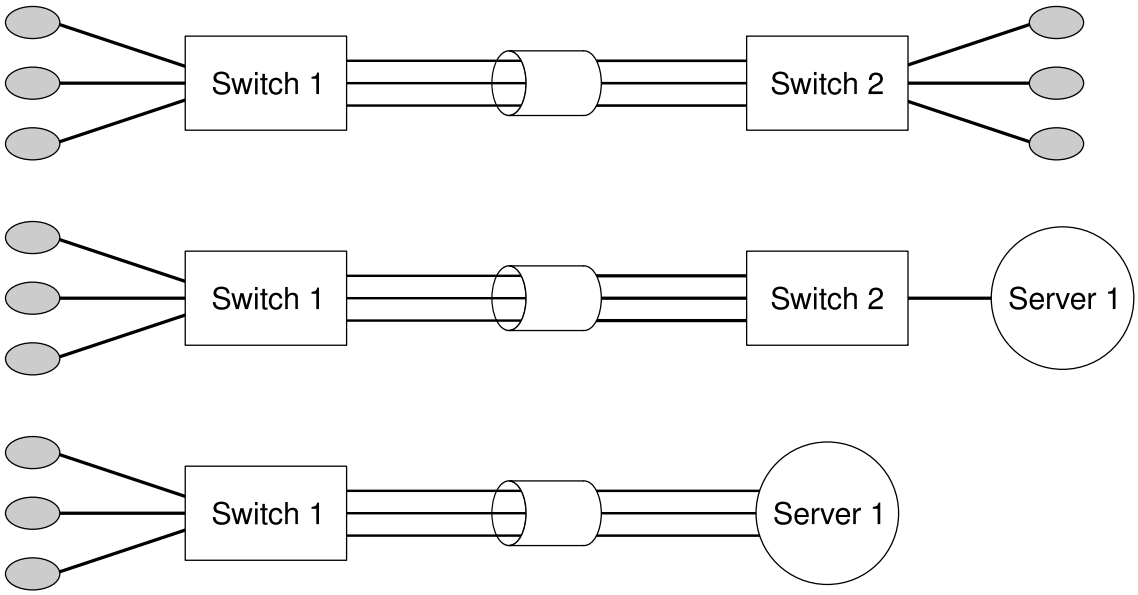

Figure 20.1. Link Aggregation Examples

20.1.1. Link Aggregation Rules

- Any port can belong to only one port trunk at a time.
- The aggregated port with the lowest port number is called the *Port Trunk Primary Port*. Other ports in the trunk are called *Secondary Ports*.
- Layer 2 features, such as STP, VLAN, CoS, and Multicast Filtering, treat a port trunk as a single link.
  - If STP puts an aggregated port in blocking or forwarding, it does so for the whole port trunk.

  - If one of the aggregated ports joins or leaves a multicast group (for example, via IGMP or GMRP), all other ports in the trunk also join or leave.
  - Any port configuration parameter changes, such as VLAN or CoS, are automatically applied to all ports in the trunk.
  - Secondary port configuration and status parameters are not be shown in configuration and status sessions in the user interface. Secondary port numbers are simply listed next to the primary port number in the user interface.
  - When a secondary port is added to a port trunk, it inherits all of the configuration settings of the primary port. When the secondary port is removed from the port trunk, the settings it had prior to the aggregation are restored.
- Physical layer features, such as physical link configuration, link status, rate limiting, and Ethernet statistics, still treat each aggregated port separately.
  - Physical configuration and status parameters are **not** automatically applied to other ports in the trunk and are displayed for each port as usual.
  - Ensure that only ports with the same speed and duplex settings are aggregated. If auto-negotiation is used, ensure that it is resolved to the same speed for all ports in the port trunk.
  - To determine the value of the Ethernet statistics counter for the port trunk, add the values of the counters for all of the ports in the port trunk.

20.1.2. Link Aggregation Limitations

- A port mirroring target port cannot be a member of a port trunk. However, a port mirroring source port can be a member of a port trunk.
- A DHCP Relay Agent Client port cannot be a member of a port trunk.
- Load balancing between the links of a bundle is randomized and may not be ideal. For example, if three 100Mbs links are aggregated, the resulting bandwidth of the port trunk may not be precisely 300Mbs.
- A Static MAC Address should not be configured to reside on an aggregated port, as it may cause some frames destined for that address to be dropped.
- A secure port cannot be a member of a port trunk.

The port trunk must be properly configured on both sides of the aggregated link. In switch-to-switch connections, if the configuration of both sides does not match (for example, some ports are mistakenly not included in the port trunk), the configuration results in a loop. The following procedure is strongly recommended to configure a port trunk:

1. *Disconnect or disable all the ports involved in the configuration. That is, disconnect or disable all ports that are being added to or removed from the port trunk.*
2. *Configure the port trunk on both switches.*
3. *Double-check the port trunk configuration on both switches.*
4. *Reconnect or re-enable the ports.*

If the port trunk is configured while the ports are not disconnected or disabled, the port will be automatically disabled for a few seconds.

The IEEE 802.3ad Link Aggregation standard requires all physical links in the port trunk to run at the same speed and in full-duplex mode. If this requirement is violated, the performance of the port trunk will drop.

If a speed/duplex mismatch is detected, the switch raises an alarm.

RSTP dynamically calculates the path cost of the port trunk based on its aggregated bandwidth. However, if the aggregated ports are running at different speeds, the path cost may not be calculated correctly.

Enabling RSTP is the best way for handling link redundancy in switch-to-switch connections composed of more than one physical link. If RSTP is enabled and increased bandwidth is not required, Link Aggregation should not be used because it may lead to a longer fail-over time.

20.2. Link Aggregation Configuration

To display the Link Aggregation menu, navigate to *interface/trunks/{trunk id}*. If no trunks are configured, see Section 20.2.1, “Configuring Port Trunks” for details on how to add trunks.

Figure 20.2. Link Aggregation menu

20.2.1. Configuring Port Trunks

Port trunks can be added by following these steps. To add ports, first go to the interface/trunks screen and enter the Edit Private mode. Click on "Add trunks".

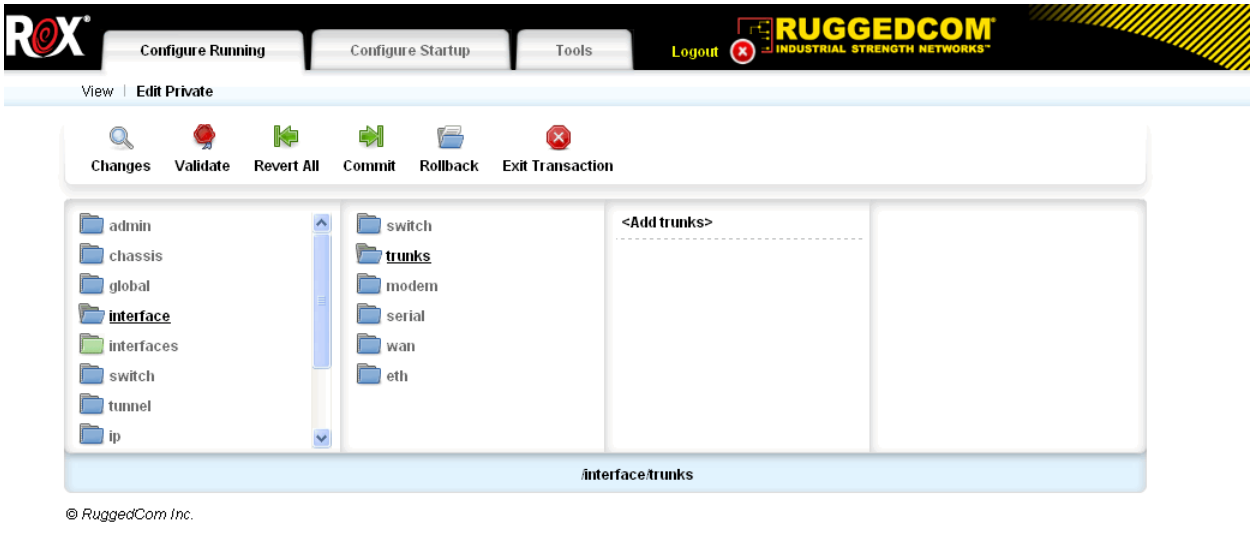

Figure 20.3. Adding Trunks

The system will now prompt you to enter a trunk ID (for example, 1) in the Key Settings form.

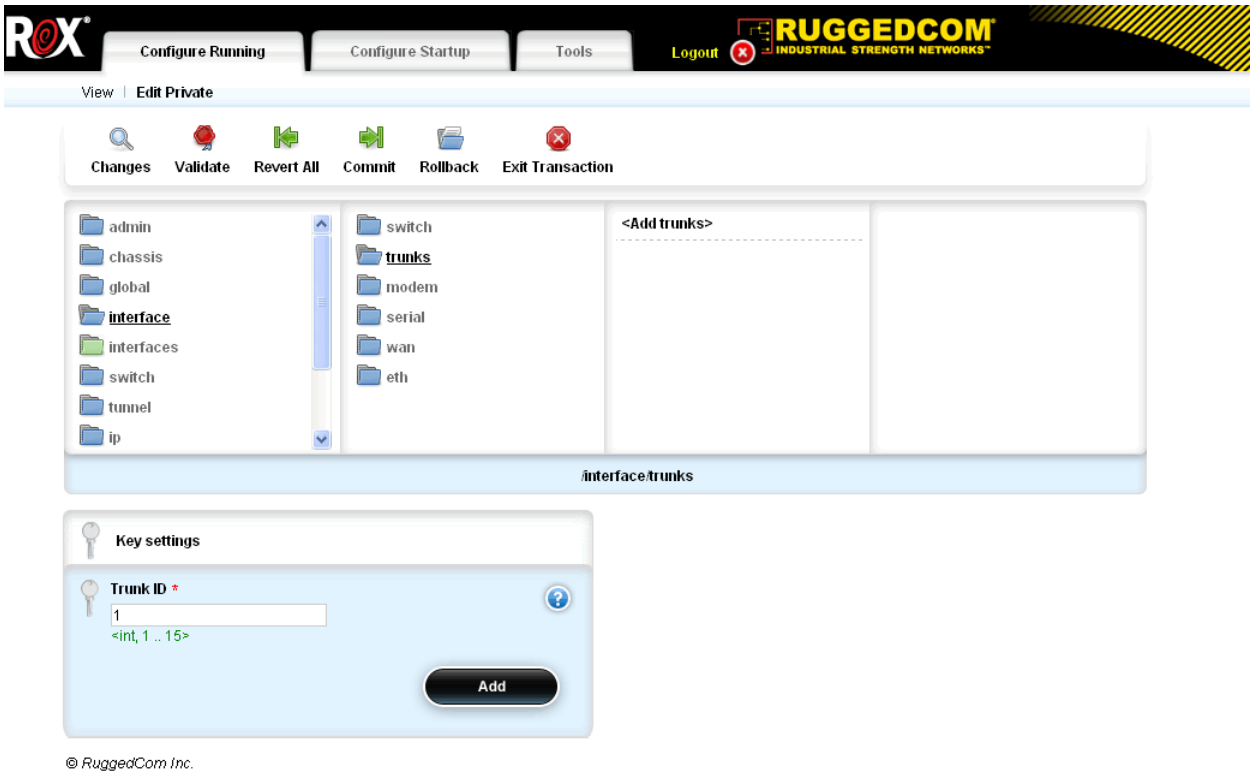

Figure 20.4. Entering a Trunk ID

Next, add parameters to the Multicast Filtering, CoS and VLAN forms.

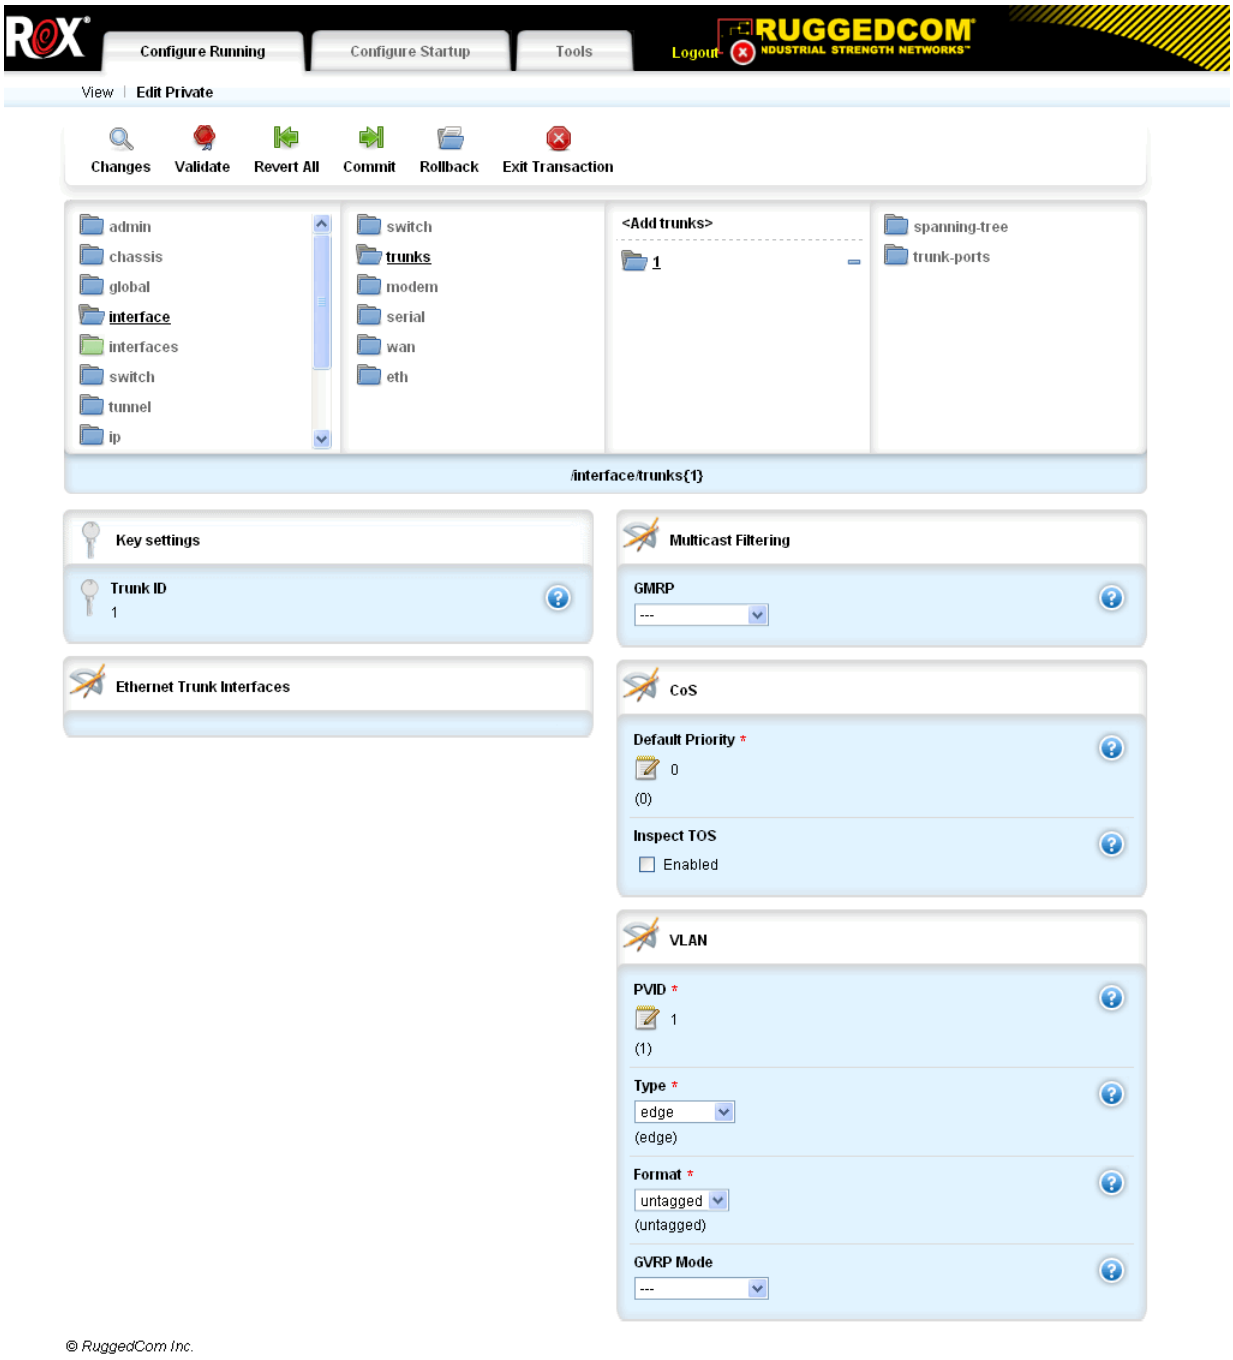

Figure 20.5. Entering Parameters for Forms

Finally, add parameters for the trunk ports. First, click on "trunk-ports" on the menu. Next, click on "Add trunk-ports" on the menu.

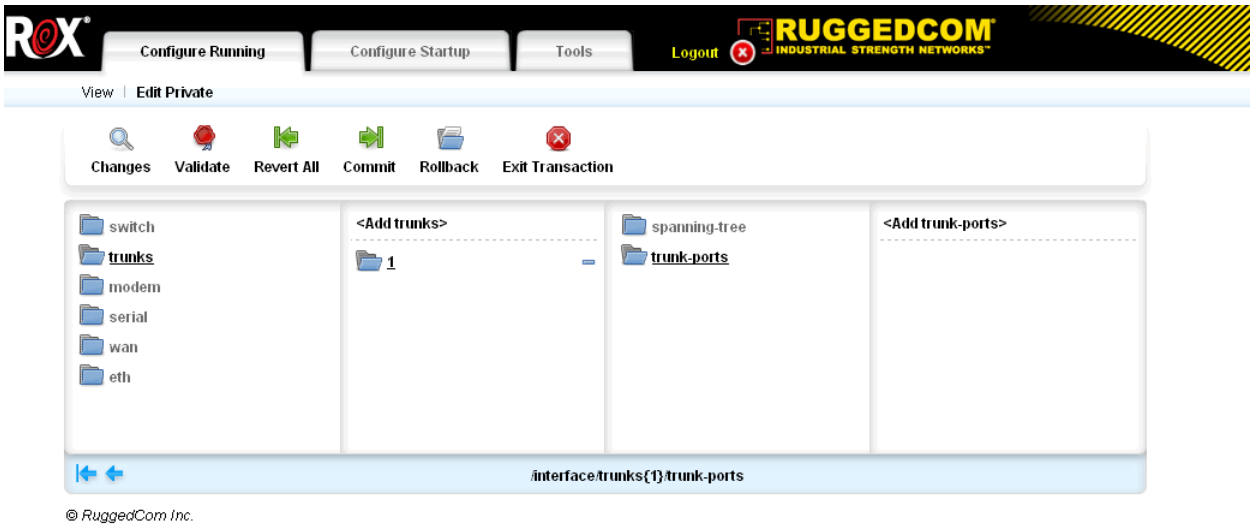

Figure 20.6. Trunk-Ports Submenu - Adding a Trunk-Port

Next, select the trunk slot from the drop-down menu on the Key Settings form. Click on "Add trunk-ports" again to add a second trunk-port. Click **Commit**. Click **Exit Transaction** when done.

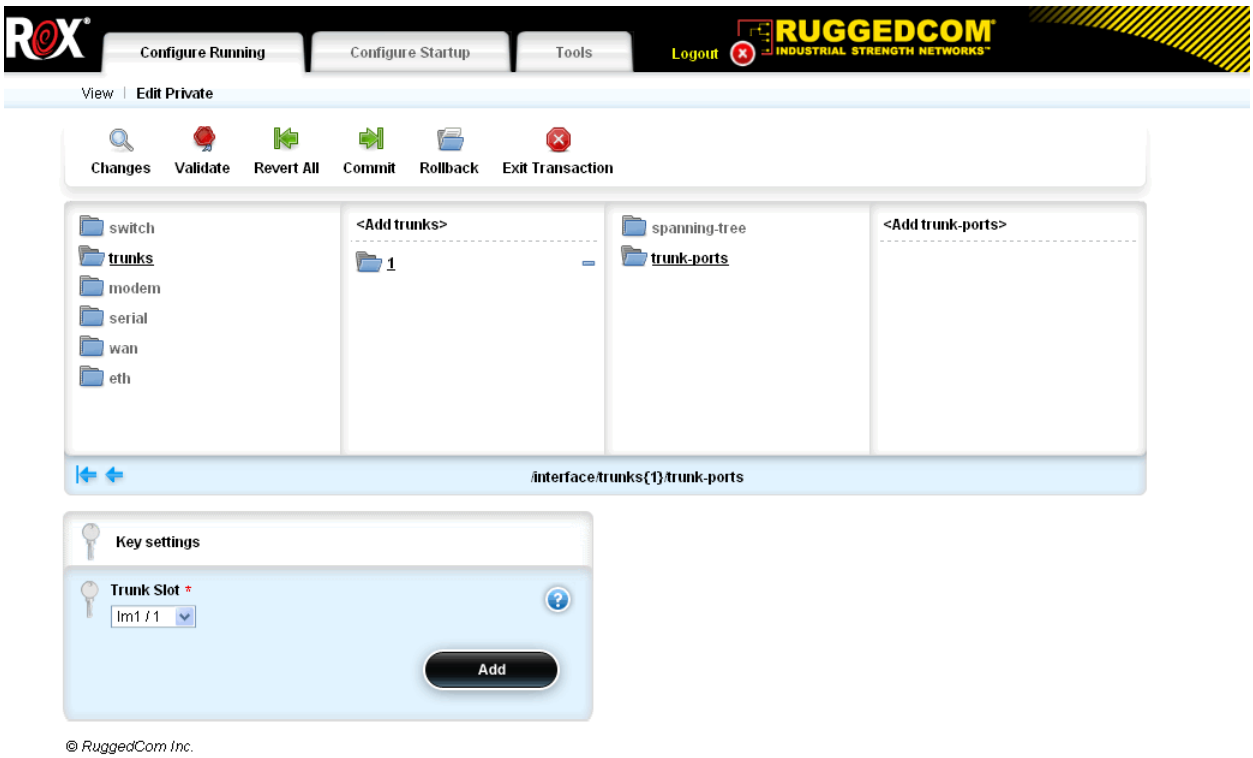

Figure 20.7. Selecting a Trunk Slot

After configuration, the Trunk Ports table (accessible at interface/trunks/{number}/trunk-ports) will display the trunk slot and trunk port. More trunk-ports can also be added by entering Edit Private mode and clicking the **Add** button that will appear in the Trunk Ports table.

Figure 20.8. Trunk Ports table

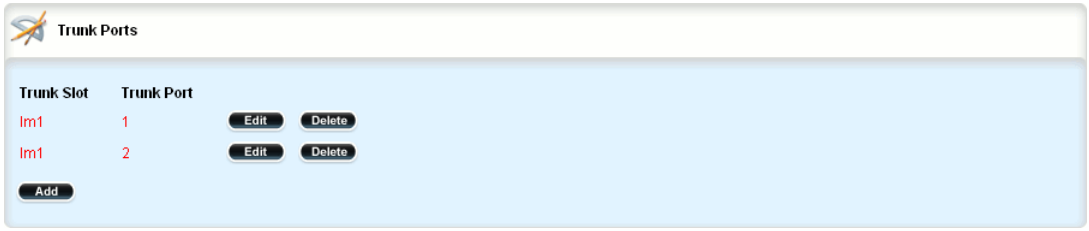

Figure 20.9. Trunk Ports Table in Edit Private Mode

To display the forms and tables below, click on *interface/trunks/{number}*. Most can also be accessed by clicking on *interface/switch/{line module}*.

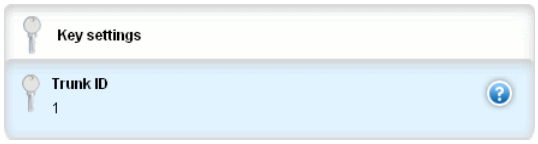

Figure 20.10. Key Settings

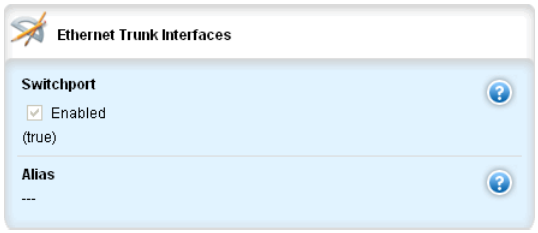

Figure 20.11. Ethernet Trunk Interfaces form

Trunk ID

Synopsis: *integer*

The trunk number. It doesn't affect port trunk operation in any way and is only used for identification.

Switchport

Synopsis: *boolean*

Default: *true*

The physical port into either Switched mode or a dedicated Routing mode.

alias

Synopsis: *A string*

The SNMP alias name of the interface

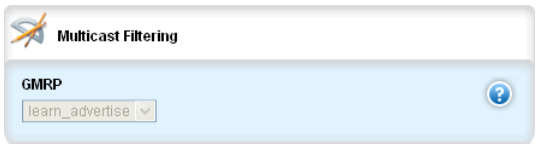

Figure 20.12. Multicast Filtering form

GMRP

Synopsis: *string - one of the following keywords { learn_advertise, advertise_only }*

GMRP (GARP Multicast Registration Protocol) operation on the port. There are several GMRP operation modes:

- DISABLED : the port is not capable of any GMRP processing.
- ADVERTISE ONLY : the port will declare all MCAST addresses existing in the switch (configured or learned) but will not learn any MCAST addresses.
- ADVERTISE and LEARN : the port will declare all MCAST Addresses existing in the switch (configured or learned) and can dynamically learn MCAST addresses.

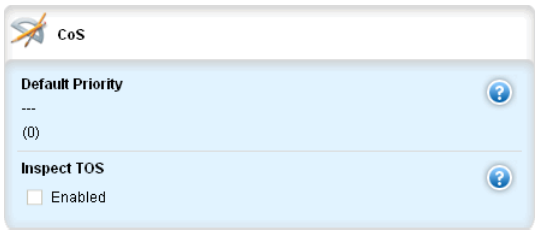

Figure 20.13. CoS form

Default Priority

synopsis: integer in the range [0 to 7]

This parameter allows to prioritize frames received on this port that are not prioritized based on the frames contents (e.g. priority field in the VLAN tag, DiffServ field in the IP header, prioritized MAC address).

Inspect TOS

This parameter enables or disables parsing of the Type-Of-Service (TOS) field in the IP header of the received frames to determine what Class of Service they should be assigned. When TOS parsing is enabled the switch will use the Differentiated Services bits in the TOS field.

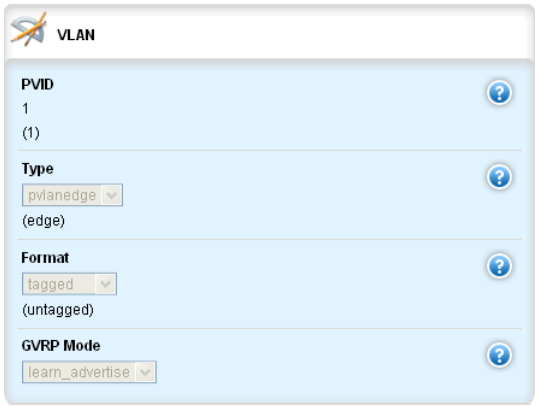

Figure 20.14. VLAN form

PVID

synopsis: integer in the range [1 to 4094]

default: 1

The Port VLAN Identifier specifies the VLAN ID associated with untagged (and 802.1p priority tagged) frames received on this port. Frames tagged with a non-zero VLAN ID will always be associated with the VLAN ID retrieved from the frame tag. Modify this parameter with care! By default, the switch is programmed to use VLAN 1 for management and every port on the switch is

programmed to use VLAN 1. If you modify a switch port to use a VLAN other than the management VLAN, devices on that port will not be able to manage the switch.

Type

synopsis: token - one of { edge, trunk, pvlanedge }

default: edge

This parameter specifies how the port determines its membership in VLANs. There are few types of ports:

EDGE - the port is only a member of one VLAN (its native VLAN specified by the 'PVID' parameter).

PVLANEdge - the port does not forward traffic to other PVLANedge ports within the same VLAN.

TRUNK - the port is automatically a member of all configured VLANs. Frames transmitted out of the port on all VLANs except the port's native VLAN will be always tagged. It can also be configured to use GVRP for automatic VLANconfiguration.

Format

synopsis: token - one of { untagged, tagged }

default: untagged

Specifies whether frames transmitted out of the port on its native VLAN(specified by the 'PVID' parameter) will be tagged or untagged.

GVRP Mode

synopsis: token - one of { advertise_only, learn_advertise }

Configures GVRP (Generic VLAN Registration Protocol) operation on the port. There are several GVRP operation modes:

DISABLED - the port is not capable of any GVRP processing.

ADVERTISE ONLY - the port will declare all VLANs existing in the switch (configured or learned) but will not learn any VLANs.

ADVERTISE and LEARN - the port will declare all VLANs existing in the switch (configured or learned) and can dynamically learn VLANs.

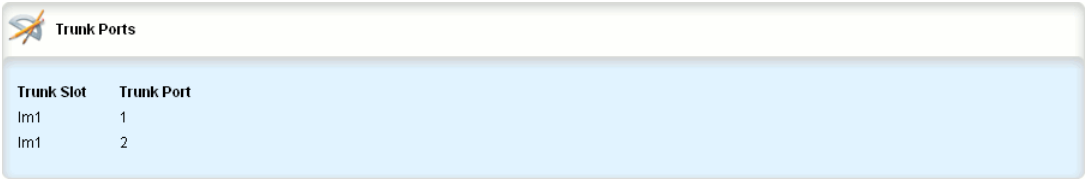

Figure 20.15. Trunk Ports table

This table displays a list of ports aggregated into the trunk. To display this table, navigate to *interface/trunks/{trunk id}/trunk-ports*.

Trunk Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Trunk Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

21. Modem

21.1. PPP and the Cellular Modem

21.1.1. PPP and Cellular Modem Fundamentals

RX1500 may be equipped with an internal cellular modem or land-line modem. PPP (the Point-to-Point Protocol) is used to establish an IP network connection over a cellular radio modem link.

Depending on local cellular network availability, one of three cellular modem types may be ordered:

- Edge
- CDMA
- HSPA

21.1.1.1. PPP Interface

When a PPP connection is established, a network interface is created in the system. You can find the interface name . Refer to the interface name when configuring firewall rules.

21.1.1.2. Authentication, IP Addressing and DNS Servers

In contrast to the configuration for land-line modems, username and password might not be required for some cellular data service providers. If username and password is not required, you can enter none in the username and password fields of the GUI, or leave them blank. If authentication is required by the cellular data service provider, again PPP authentication will automatically use PAP or CHAP. Your service provider will provide you with a username and password along with an Access Provider Name (APN), which must be entered in the GUI.

The authentication process will provide a local IP address for use on the PPP interface and optionally the addresses of the DNS servers and a default gateway address to use. You should generally use these addresses unless you need to provide your own.

The PPP interface's IP address, obtained from the PPP server, can be either a dynamic or a static IP address. Firewall configuration should be performed as is appropriate.

21.1.1.3. When the Modem Connects

A PPP Client Connection for the cellular modem may be configured to connect at boot time.

21.1.2. PPP Cellular Modem Information and Configuration

The following sections review the forms used to view and configure HSPA, Edge, and CDMA cellular modems. The HSPA, Edge, and CDMA menus can be accessed from the *interfaces/cellmodem* menu below.

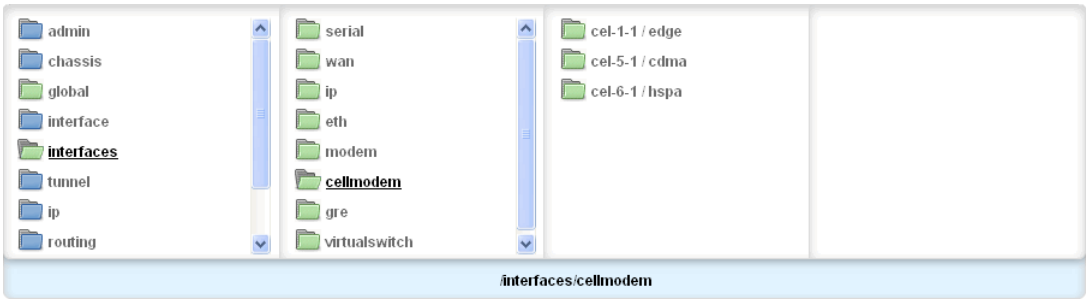

Figure 21.1. Interfaces Cellmodem menu

21.1.2.1. HSPA

The HSPA GSM profile is selected from the HSPA menu but Edge data needs to be configured from the Global Cellular GSM menu. See Section 21.1.2.3, “Global Cellular Modem GSM Configuration” for information on configuration.

If data is configured, the HSPA Cellular Modem Information form can be found under *interfaces/cellmodem/{line module}/hspa*.

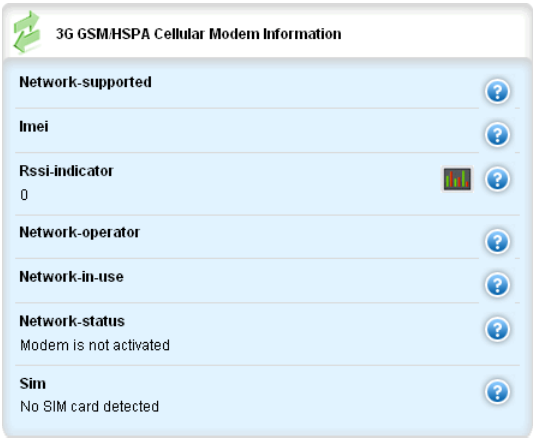

Figure 21.2. HSPA Cellular Modem Information form

The HSPA Cellular Modem Information form displays modem information.

network-supported

Synopsis: *A string*

Wireless technologies supported by the modem

imei

Synopsis: *A string*

International Mobile Equipment Indentity

radio

Synopsis: *A string*

The current RF status of cellmodem

rssi-indicator

Synopsis: *integer*

The Received Signal Strength Indicator in dBm

network-operator

Synopsis: *A string*

The wireless network operator currently in use

network-in-use

Synopsis: *A string*

The network technology currently in use by the modem

network-status

Synopsis: *A string*

The registration status of the modem with the wireless network

sim

Synopsis: *A string*

The Subscriber Indentity Module number

The following information provides additional details about the fields in the HSPA Cellular Modem Information Form.

The *IMEI* (International Mobile Equipment Identity) is a numeric identifier unique to the cellular modem card.

Rssi Indicator (Received Signal Strength) indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

The *Network In Use* field displays which network technology is currently in use between the modem and the network.

Network Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

SIM displays the ID of the SIM card currently installed in the cellular modem.

21.1.2.2. Edge

The Edge GSM profile is selected from the Edge menu but Edge data needs to be configured from the Global Cellular GSM menu. See Section 21.1.2.3, “Global Cellular Modem GSM Configuration” for information on configuration.

If data is configured, the path to the Edge Cellular Modem Information form is *interfaces/cellmodem/{line module}/edge*.

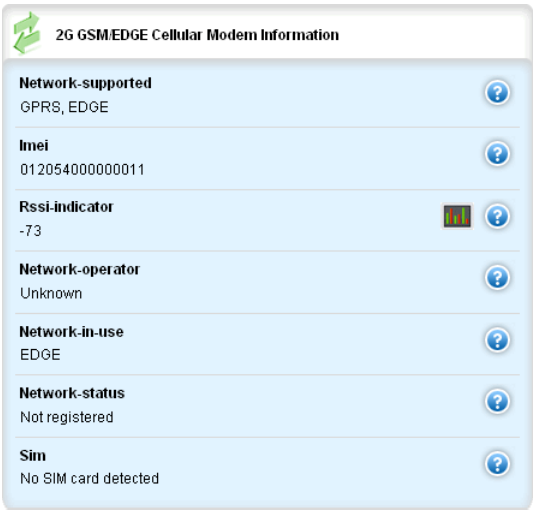

Figure 21.3. Edge Cellular Modem Information form

network-supported

Synopsis: *A string*

Wireless technologies supported by the modem

imei

Synopsis: *A string*

International Mobile Equipment Indentity

radio

Synopsis: *A string*

The current RF status of cellmodem

rssi-indicator

Synopsis: *integer*

The Received Signal Strength Indicator in dBm

network-operator

Synopsis: *A string*

The wireless network operator currently in use

network-in-use

Synopsis: *A string*

The network technology currently in use by the modem

network-status

Synopsis: *A string*

The registration status of the modem with the wireless network

sim

Synopsis: *A string*

The Subscriber Indentity Module number

The following information provides additional details about the fields in the Edge Cellular Modem Information Form.

Rssi Indicator (Received Signal Strength) indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

Network Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

SIM displays the ID of the SIM card currently installed in the cellular modem.

21.1.2.3. Global Cellular Modem GSM Configuration

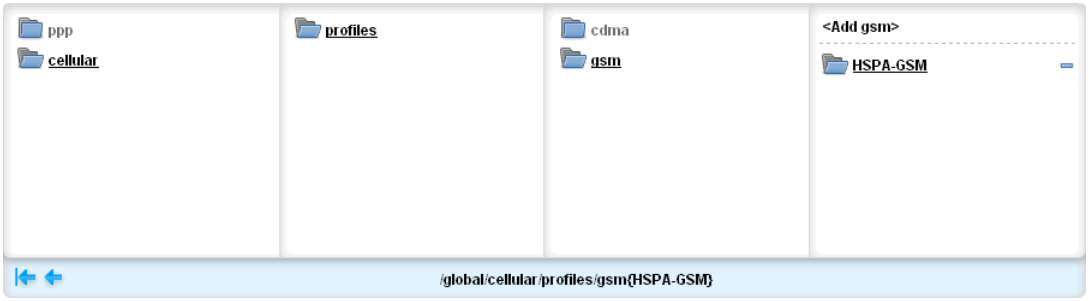

Figure 21.4. Global Cellular GSM menu

HSPA and Edge data is configured from the Global Cellular GSM menu.

Figure 21.5. GSM Cellular Network Configuration form

name

Synopsis: *A string*

Create gsm profile name

apn

Synopsis: *string*

The wireless network access point name

dial-string

Synopsis: *string*

Default: **99***1#*

The dial string given by the wireless provider to connect to the access point name

The Access Point Name (APN) is necessary only on GPRS networks (Edge or HSPA). It is the name of the cellular network access point which provides a gateway to the Internet. This information will be provided by the wireless network when you register for data service. This field is not used for CDMA modems.

The Dial string is a special command to be sent by the cellular modem to the cellular network to establish a data connection. For example, for GSM/GPRS networks, this is typically: *99***1#. This command will depend on the wireless network. Please consult the wireless network operator for the correct dial string command for data service. A regular telephone number is usually not required to connect to a GSM/GPRS network.

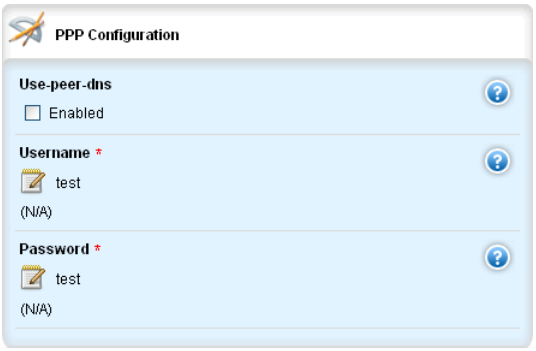

Figure 21.6. PPP Configuration form

use-peer-dns

Enables the DNS server entries that the PPP server recommends. Enables this option unless you provide your own name servers

username

Synopsis: *string*

Default: *N/A*

The user ID to connect to the remote server

password

Synopsis: *string*

Default: *N/A*

The password to be authenticated by the remote server

dial-on-demand

Activates Dial-on-Demand on this connection. The establishment of the PPP connection is postponed until there is data to be transmitted via the interface

disconnect-idle-timeout

Synopsis: *integer*

Default:

The time in seconds to wait before disconnecting PPP when there is no traffic on the link. This option is only valid when Dial-on-Demand is enabled

failover-on-demand

Activate link failover on-demand on this device. PPP link establishment on this device is controlled by link failover

21.1.2.4. CDMA

The CDMA GSM profile is selected by using the CDMA EVDO Cellular Modem Configuration form but the profile needs to be configured first from the Global Cellular CDMA Modem menu. See Section 21.1.2.4.2, “Global Cellular CDMA Modem Configuration” for information on configuration. After the profile has been configured, the CDMA functions and the CDMA EVDO modem Configuration form are accessible from the CDMA menu.

The path to the CDMA menu is *interfaces/cellmodem/{line module}/cdma*. The CDMA EVDO Cellular Modem Information form appears on the same screen as the *CDMA* menu.

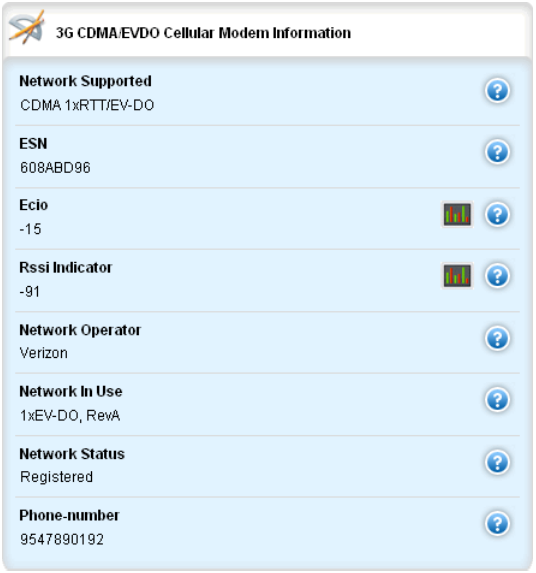

Figure 21.7. CDMA EVDO Cellular Modem Information form

network-supported

Synopsis: *A string*

Wireless technologies supported by the modem

esn

Synopsis: *A string*

The Electronic Serial Number of the modem. ESN is only avaible for the CDMA modem.

ecio

Synopsis: *integer*

The total energy per chip per power density value in dBm

rssi-indicator

Synopsis: *integer*

The Received Signal Strength Indicator in dBm

network-operator

Synopsis: *A string*

The wireless network operator currently in use

network-in-use

Synopsis: *A string*

The network technology currently in use by the modem

network-status

Synopsis: *A string*

The registration status of the modem with the wireless network

phone-number

Synopsis: *A string*

The subscriber phone number of the CDMA modem

The information below provides additional details about the fields in the CDMA EVDO Cellular Modem Information form.

The *Electronic Serial Number (ESN)* is a numeric identifier unique to the cellular modem card. This corresponds to the IMEI for GSM networks.

Rssi Indicator (Received Signal Strength) indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

The *Network In Use* field displays which network technology is currently in use between the modem and the network.

Network Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

Phone Number displays the cellular telephone number associated with the account created to provide service for the modem.

21.1.2.4.1. Cellular Modem Account Activation

Prior to use, a CDMA-type cellular modem must be activated for use on a particular provider's network. Once the activation process has been completed, the modem will be able to connect to the network without further intervention. Two account activation methods are provided by ROX™: "OTA (Over-the-Air)" and "Manual". Both activation methods are described in this section.

Over-The-Air Account Activation

ROX™ supports the OTASP (Over-the-Air Service Provisioning) mechanism offered by most CDMA cellular service providers for provisioning cellular end stations for use on their networks. Using this method, the service provider, or carrier, supplies an OTASP dial string which ROX™ can use to contact the cellular network via the modem. During this OTASP call, the carrier authorizes and configures the modem for use on its network. Note that an OTASP dial string typically begins with "*228".

The path to the CDMA Over The Air Activation form and Trigger Action form is *interface/modem/lm6 / 1/cdma/OverTheAirActivation*.

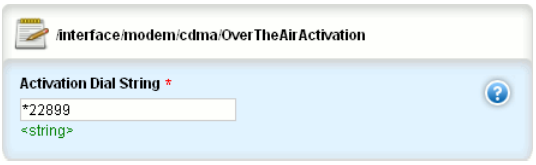

Figure 21.8. CDMA Over The Air Activation form

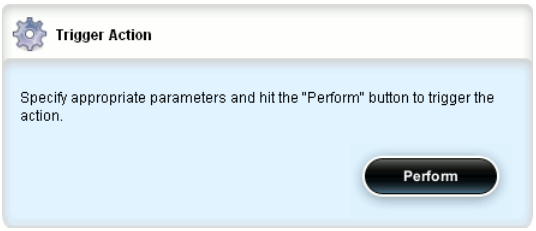

Figure 21.9. CDMA Over The Air Activation Trigger Action form

1. First, establish an account with the help of a service representative of the cellular network provider.
2. Enter the OTASP dial string supplied in the *Activation Dial string* field of the Over The Air Activation form.
3. Click the *Perform* button on the *Trigger Action* form.

Manual Account Activation

If the carrier does not support Over-the-Air Service Provisioning, the cellular modem must be programmed via the *Manual Account Activation* form using settings supplied by the carrier's service personnel:

The path to the CDMA Manual Activation form and Trigger Action form is *interface/modem/lm6 / 1/ cdma/ManualActivation*.

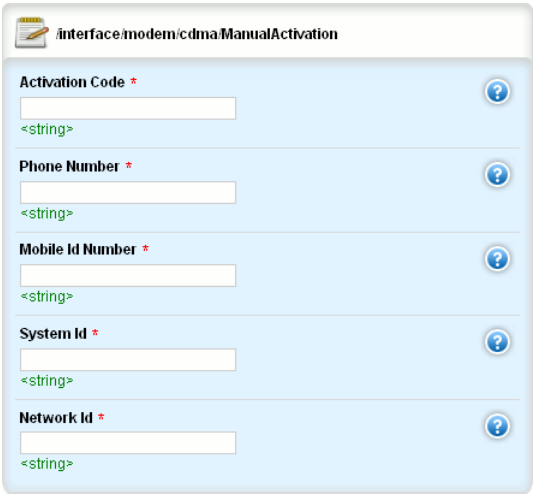

Figure 21.10. CDMA Manual Activation form

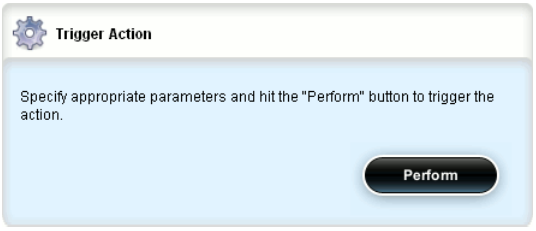

Figure 21.11. CDMA Manual Activation Trigger Action form

1. First, establish an account with a service representative of the cellular network provider. You will need the following settings in order to activate your modem. Note that not all of these parameters are required by all network providers:
   - *Activation code*, also known as a "subsidy lock".
   - *Phone Number*, or MDN (Mobile Directory Number).
   - *MIN* (Mobile Identification Number), often the same as the Phone Number.
   - *System ID*, or Home System ID.
   - *Network ID*
2. After specifying the parameters above in the Manual Activation form, click the *Perform* button on the *Trigger Action* form.

Reset Modem

The modem can be reset using the Reset Modem Trigger Action form. The path to the CDMA Reset Modem Trigger Action form is *interface/modem/lm6 / 1/cdma/resetmodem*.

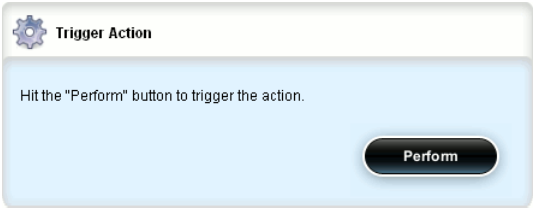

Figure 21.12. CDMA Reset Modem Trigger Action form

21.1.2.4.2. Global Cellular CDMA Modem Configuration

Figure 21.13. Global Cellular CDMA menu

The path to this menu is *global/cellular/profiles/cdma*. The Cellular Network Configuration table appears on the same screen as the *global* menu.

CDMA data is configured from the Global Cellular CDMA menu.

Figure 21.14. Cellular Network Configuration table

Figure 21.15. Cellular Network Configuration form

The Cellular Network Configuration form and the PPP Configuration form appear on the same screen as the *global* menu.

name

Synopsis: *A string*

Create cdma profile name

dial-string

Synopsis: *A string*

Default: *#777*

The dial string to connect to the wireless provider

The Dial string is a special command to be sent by the cellular modem to the cellular network to establish a data connection. For example, for GSM/GPRS networks, this is typically: *99***1#. This command will depend on the wireless network. Please consult the wireless network operator for the correct dial string command for data service. A regular telephone number is usually not required to connect to a GSM/GPRS network.

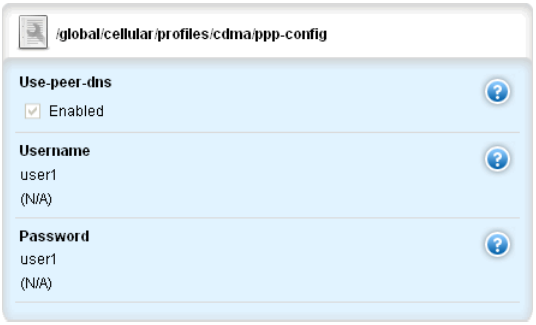

Figure 21.16. PPP Configuration form

use-peer-dns

Enables the DNS server entries that the PPP server recommends. Enables this option unless you provide your own name servers

username

Synopsis: *string*

Default: *N/A*

The user ID to connect to the remote server

password

Synopsis: *string*

Default: *N/A*

The password to be authenticated by the remote server

dial-on-demand

Activates Dial-on-Demand on this connection. The establishment of the PPP connection is postponed until there is data to be transmitted via the interface

disconnect-idle-timeout

Synopsis: *integer*

Default:

The time in seconds to wait before disconnecting PPP when there is no traffic on the link. This option is only valid when Dial-on-Demand is enabled

failover-on-demand

Activate link failover on-demand on this device. PPP link establishment on this device is controlled by link failover

21.1.2.5. CellModem

21.1.2.5.1. CellModem Configuration

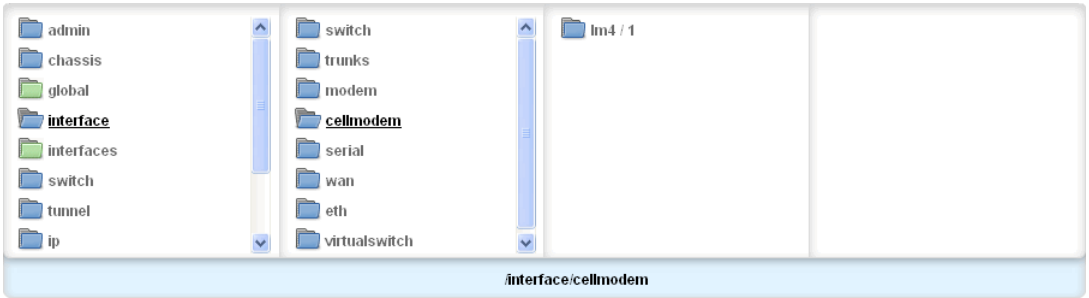

Figure 21.17. Interface Cellmodem menu

The path to the *interface/cellmodem* menu is *interface/cellmodem*. The Routable Cellular Modem Interfaces table appears on the same screen as this menu.

Figure 21.18. Routable Cellular Modem Interfaces table

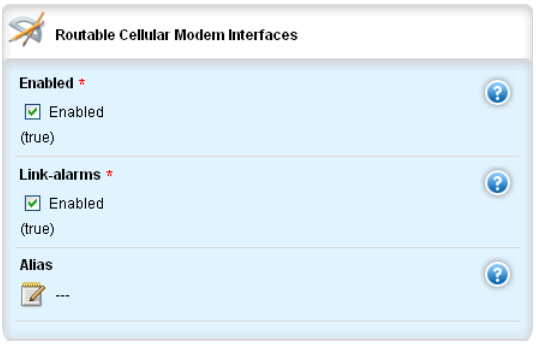

Figure 21.19. Routable Cellular Modem Interfaces form

The path to this form is *interface/cellmodem/{line module}*.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

enabled

Synopsis: *boolean*

Default: *true*

Provides the option to enable or disable this interface. When unchecked(i.e disabled), the interface will prevent all frames from being sent and received on that interface.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

alias

Synopsis: *A string*

The SNMP alias name of the interface

Figure 21.20. Interface Cellmodem HSPA menu

The path to this menu is *interface/cellmodem/{line module}/hspa*.

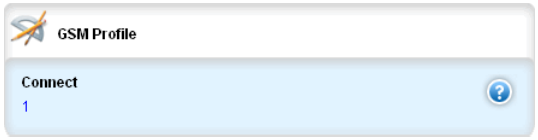

Figure 21.21. GSM Profile form

The path to this form is *interface/cellmodem/{line module}/hspa/ppp-client*.

Connect

Synopsis: *A string*

Selects the gsm profile to connect to wireless network. The gsm profile is configured in /global/cellular/profiles/gsm

21.1.2.5.2. CellModem Status

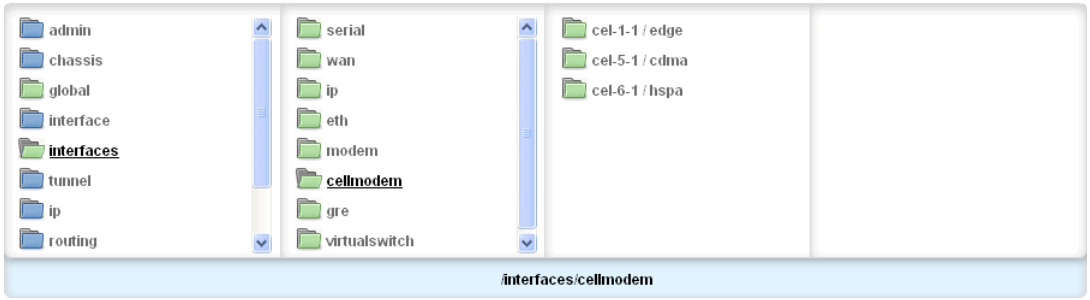

Figure 21.22. Interfaces Cellmodem menu

The path to the *interfaces/cellmodem* menu is *interfaces/cellmodem*. The Cellular Modem Interfaces table appears on the same screen as this menu.

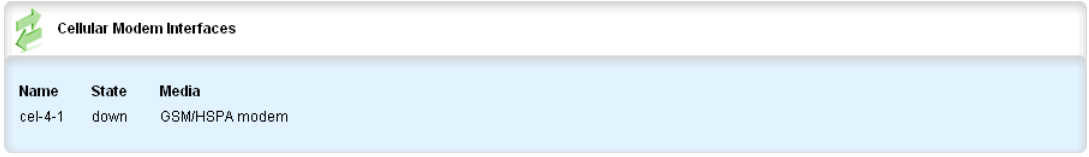

Figure 21.23. Cellular Modem Interfaces table

Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

Interface name

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

The port's link status.

media

Synopsis: *A string*

The type of port media { ***range of values** }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

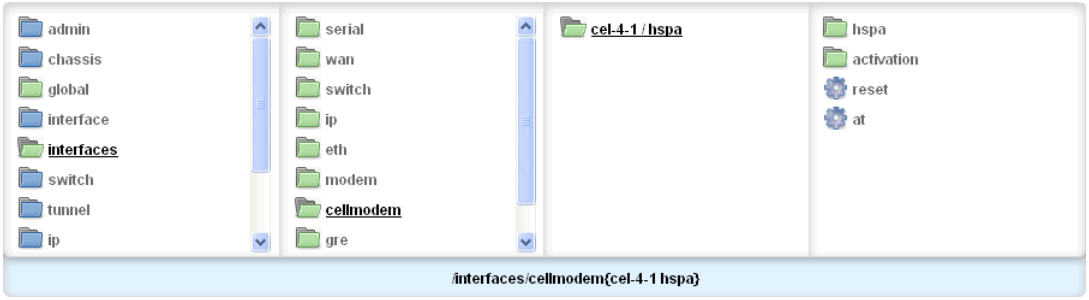

Figure 21.24. Interfaces Cellmodem HSPA menu

The path to this menu is *interfaces/cellmodem/{line module}/hspa*. The Cellular Modem Interfaces form and the HSPA PPP Interfaces Statistics form appear on the same screen as this menu.

Figure 21.25. Cellular Modem Interfaces form

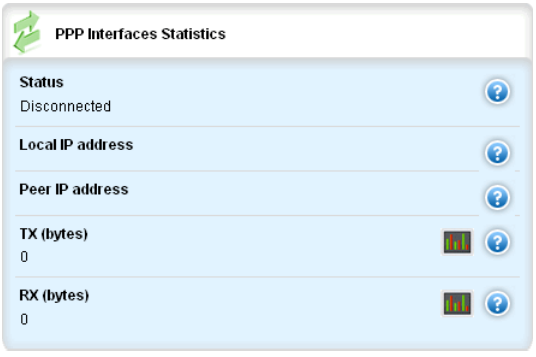

Figure 21.26. HSPA PPP Interfaces Statistics form

Status

Synopsis: *A string*

PPP connection status

Local IP address

Synopsis: *A string*

The IP address assigned to the modem by the remote server

Peer IP address

Synopsis: *A string*

The IP address of the remote server

TX (bytes)

Synopsis: *unsigned integer*

The bytes transmitted over the modem

RX (bytes)

Synopsis: *unsigned integer*

The bytes received by the modem

22. Serial Protocols

22.1. Introduction

This chapter familiarizes the user with:

- RawSocket Applications
- TCP Modbus Server Applications
- DNP (Distributed Network Protocol)
- Configuring Serial ports for RawSocket
- Viewing Serial Port and TCP Connection Status and Statistics
- Resetting Serial ports

22.1.1. Serial IP Port Features

The RuggedCom Serial Server provides the following features for forwarding serial traffic over IP:

- Raw Socket Protocol - a means to transport streams of characters from one serial port on the router to a specified remote IP address and TCP port
- RX1500 supports up to 24 serial ports; RX1501 supports up to 36 serial ports
- Bit rates of 300, 600, 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200, or 230400 bps.
- Supports RS232, RS422 and RS485 party line operation.
- XON/XOFF flow control.
- Supports a point-to-point connection mode and a broadcast connection mode in which up to 32 remote servers may connect to a central server.
- TCP/IP incoming, outgoing or both incoming/outgoing connections mode, configurable local and remote TCP port numbers.
- Packetize and send data on a full packet, a specific character or upon a timeout.
- Supports a “turnaround” time to enforce minimum times between messages sent out the serial port.
- Debugging facilities including connection tracing and statistics

22.1.1.1. LED Designations

Each RJ45 connector on the 6S01 serial module features a single LED. The LED illuminates to indicate serial traffic.

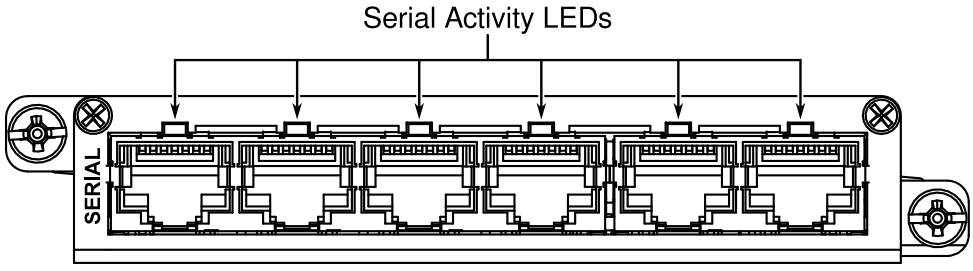

Figure 22.1. 6S01 Serial Module RJ45 Connector LEDs

22.1.2. Serial Protocols Applications

22.1.2.1. Character Encapsulation

Character encapsulation is used any time a stream of characters must be reliably transported across a network.

The character streams can be created by any serial device. The baud rates supported at either server need not be the same. If configured, the router will obey XON/XOFF flow control from the end devices.

One of the servers is configured to listen to TCP connection requests on a specific TCP port number. The other server is configured to connect to its peer on the listening port number. The RX1500 will periodically attempt to connect if the first attempt fails and after a connection is broken.

The RX1500 can be used to connect to any device supporting TCP (e.g. a host computer's TCP stack or a serial application on a host using port redirection software).

22.1.2.2. RTU Polling

The following applies to a variety of RTU protocols besides ModBus RTU, including ModBus ASCII and DNP.

The remote router communicates with host equipment through:

- native TCP connections,
- another RX1500 via a serial port or
- a port redirection package which Supports TCP.

If a RX1500 is used at the host end, it will wait for a request from the host, encapsulate it in a TCP message, and send it to the remote side. There, the remote RX1500 will forward the original request to the RTU. When the RTU replies, the the RX1500 will forward the encapsulated reply back to the host end.

ModBus does not employ flow-control so XON/XOFF should not be configured.

The RX1500 maintains configurable timers to help decide replies and requests are complete and to handle special messages such as broadcasts.

The RX1500 will also handle the process of line-turnaround when used with RS485.

22.1.2.3. Broadcast RTU Polling

Broadcast polling allows a single host connected RX1500 to "fan-out" a polling stream to a number of remote RTUs.

The host equipment connects via a serial port to a RX1500. Up to 32 remote devices may connect to the host server via the network.

Initially, the remote servers will place connections to the host server. The host server in turn is configured to accept the required number of incoming connections.

The host will sequentially poll each RTU. Each poll received by the host server is forwarded (i.e. broadcast) to all of the remote servers. All RTUs will receive the request and the appropriate RTU will issue a reply. The reply is returned to the host server, where it is forwarded to the host.

22.1.3. Serial Protocols Concepts And Issues

22.1.3.1. Host And Remote Roles

The RX1500 can either initiate or accept a TCP connection for serial encapsulation. It can establish a connection from field (“remote”) equipment to the central site (“host”) equipment, vice versa, or bi-directionally.

Configure the RX1500 at the host end to connect to the remote when:

- The host end uses a port redirector that must make the connection.
- The host end is only occasionally activated and will make the connection when it becomes active.
- A host end firewall requires the connection to be made outbound.

Connect from the remote to the host if the host end accepts multiple connections from remote ends in order to implement broadcast polling.

Connect from each side to other if both sides support this functionality.

22.1.3.2. Use Of Port Redirectors

Port redirectors are PC packages that emulate the existence of communications ports. The redirector software creates and makes available these “virtual” COM ports, providing access to the network via a TCP connection.

When a software package uses one of the virtual COM ports, a TCP connection is placed to a remote IP address and TCP port that has been programmed into the redirector. Some redirectors also offer the ability to receive connections.

22.1.3.3. Message Packetization

The server buffers received characters into packets in order to improve network efficiency and demarcate messages.

The server uses three methods to decide when to packetize and forward the buffered characters to the network:

- Packetize on Specific Character,
- Packetize on timeout and
- Packetize on full packet.

If configured to packetize on a specific character, the server will examine each received character and will packetize and forward upon receiving the specific character. The character is usually a <CR> or an <LF> character but may be any ASCII character.

If configured to packetize on a timeout, the server will wait for a configurable time after receiving a character before packetizing and forwarding. If another character arrives during the waiting interval, the timer is restarted. This method allows characters transmitted as part of an entire message to be forwarded to network in a single packet, when the timer expires after receiving the very last character of the message. This is usually the only packetizer selected when supporting ModBus communications.

Finally, the server will always packetize and forward on a full packet, i.e. when the number of characters fills its communications buffer (1024 bytes).

22.1.3.4. Use of Turnaround Delays

Some RTU protocols (such as ModBus) use the concept of a turnaround delay. When the host sends a message (such as a broadcast) that does not invoke an RTU response, it waits a turnaround delay

time. This delay ensures that the RTU has time to process the broadcast message before it has to receive the next poll.

When polling is performed, network delays may cause the broadcast and next poll to arrive at the remote server at the same time. Configuring a turnaround delay will enforce a minimum separation time between each message sent out the port. Note that turnaround delays do not need to be configured at the host computer side and may be disabled there.

22.1.4. TcpModBus Server Application

The TcpModbus Server application is used to transport Modbus requests and responses across IP networks.

The source of the polls is a Modbus “master”, a host computer that issues the polls over a serial line.

A TcpModbus Client application, such as that implemented by the RuggedServer accepts Modbus polls on a serial line from a master and determines the IP address of the corresponding RTU. The client then encapsulates the message in TCP and forwards the frame to a Server Gateway or native TcpModbus RTU. Returning responses are stripped of their TCP headers and issued to the master.

The TcpModbus Server application accepts TCP encapsulated modbus messages from Client Gateways and native masters. After removing the TCP headers the messages are issued to the RTU. Responses are TCP encapsulated and returned to the originator.

A “native” TcpModbus master is one that can encapsulate the Modbus polls in TCP and directly issue them to the network.

22.1.4.1. Local Routing At The Server Gateway

The Server Gateway supports up to 32 RTUs on any of its four ports. When a request for a specific RTU arrives the server will route it to the correct port.

22.1.4.2. MultiMaster Capability

It is possible for multiple masters to simultaneously issue requests for the same RTU. The Server Gateway will queue the requests and deliver them to the RTU in turn. This “multimaster” capability allows widely distributed masters to configure and extract information from the RTU.

22.1.5. TcpModbus Concepts And Issues

22.1.5.1. Host And Remote Roles

Client gateways (such as that implemented by the RX1500) always make the TCP connection to the Server Gateway. The Server Gateway can only accept a connection.

22.1.5.2. Port Numbers

The TCP port number dedicated to Modbus use is port 502. The Server Gateway can also be configured to accept a connection on a configurable port number. This auxiliary port can be used by masters that do not support port 502.

The Server Gateway is capable of creating only one connection on the specified auxiliary port, whereas when Modbus is configured to use the default port, 502, it may connect to multiple RTUs.

22.1.5.3. Retransmissions

The Server Gateway offers the ability to resend a request to an RTU should the RTU receive the request in error or the Server Gateway receives the RTU response in error.

The decision to use retransmissions, and the number to use depends upon factors such as:

- The probability of a line failure
- The number of RTUs and amount of traffic on the port
- The cost of retransmitting the request from the server vs. timing-out and retransmitting at the master. This cost is affected by the speed of the ports and of the network.

22.1.5.4. ModBus Exception Handling

If the Server Gateway receives a request for an unconfigured RTU, it will respond to the originator with a special message called an exception (type 10). A type 11 exception is returned by the server if the RTU fails to respond to requests.

Native TcpModbus polling packages will want to receive these messages. Immediate indication of a failure can accelerate recovery sequences and reduce the need for long timeouts.

22.1.5.5. TcpModbus Performance Determinants

The following description provides some insight into the possible sources of delay and error in an end-to-end TcpModbus exchange.

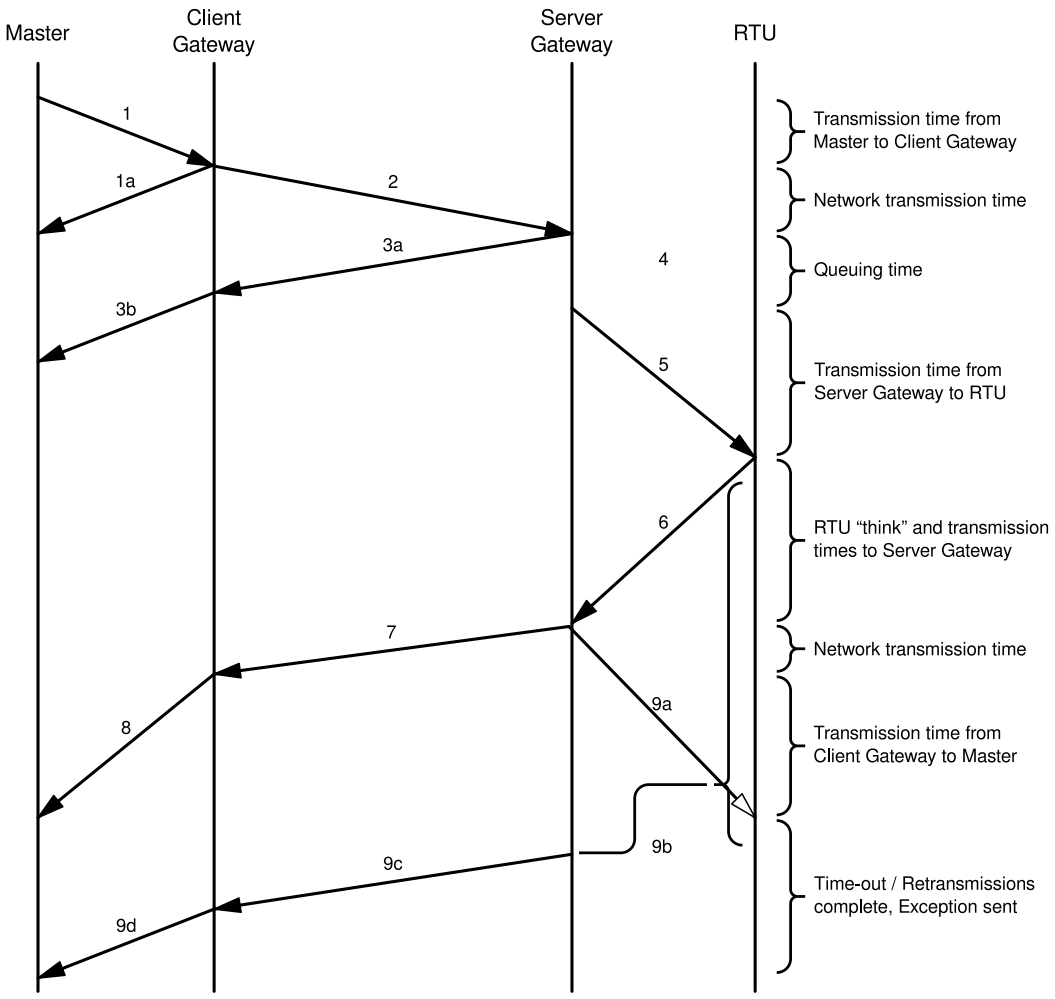

Figure 22.2. Sources of Delay and Error in an End to End Exchange

In step 1, the master issues a request to the Client Gateway. If the Client Gateway validates the message it will forward it to the network as step 2.

The Client Gateway can respond immediately in certain circumstances, as shown in step 1a. When the Client Gateway does not have a configuration for the specified RTU it will respond to the master with an exception using TcpModbus exception code 11 (“No Path”). When the Client Gateway has a configured RTU but the connection is not yet active it will respond to the master with an exception using TcpModbus exception code 10 (“No Response”). If the forwarding of TcpModbus exceptions is disabled, the client will not issue any responses.

Steps 3a and 3b represents the possibility that the Server Gateway does not have configuration for the specified RTU. The Server Gateway will always respond with a type 10 (“No Path”) in step 3a, which the client will forward in step 3b.

Step 4 represents the possibility of queuing delay. The Server Gateway may have to queue the request while it awaits the response to a previous request. The worst case occurs when a number of requests are queued for an RTU that has gone offline, especially when the server is programmed to retry the request upon failure.

Steps 5-8 represent the case where the request is responded to by the RTU and is forwarded successfully to the master. It includes the “think time” for the RTU to process the request and build the response.

Step 9a represents the possibility that the RTU is offline, the RTU receives the request in error or that the Server Gateway receives the RTU response in error. If the Server Gateway does not retry the request, it will issue an exception to the originator.

22.1.5.6. A Worked Example

A network is constructed with two Masters and 48 RTUs on four Server Gateways. Each of the Master is connected to a Client Gateway with a 115.2 Kbps line. The RTUs are restricted to 9600 bps lines. The network is Ethernet based and introduces an on average 3 ms of latency. Analysis of traces of the remote sites has determined that the min/max RTU think times were found to be 10/100 ms. What time-out should be used by the Master?

The maximum sized Modbus message is 256 bytes in length. This leads to a transmission time of about 25 ms at the Master and 250 ms at the RTU. Under ideal circumstances the maximum round trip time is given by: 25 ms (Master->client) + 3 ms (network delay) + 250 ms (server->RTU) + 100 ms (Think time) + 250 ms (RTU->server) + 3 ms (network delay) + 25 ms (client->Master). This delay totals about 650 ms.

Contrast this delay with that of a “quick” operation such as reading a single register. Both request and response are less than 10 bytes in length and complete (for this example) in 1 and 10 ms at the client and server. Assuming the RTU responds quickly, the total latency will approach 35 ms.

It is also necessary to take account such factors as the possibility of line errors and collisions between masters at the server.

The server may be configured to recover from a line error by retransmitting the request. Given a maximum frame transmission time of 250 ms and an RTU latency of 100 ms, it would be wise to budget 350 ms for each attempt to send to the RTU. Configuring a single retransmission would increase the end-to-end delay from about 650 ms to about 1000 ms.

The server can already be busy sending a request when the request of our example arrives. Using the figures from the above paragraph, the server being busy would increase the end-to-end delay from 1000 to 1350 ms.

The preceding analysis suggests that the Master should time-out at some time after 1350 ms from the start of transmission.

22.1.6. DNP (Distributed Network Protocol)

The RX1500 supports DNP 3.0, commonly used by utilities in process automation systems. DNP3 protocol messages specify source and destination addresses. A destination address specifies which device should process the data, and the source address specifies which device sent the message. Having both destination and source addresses satisfies at least one requirement for peer-to-peer communication since the receiver knows where to direct a response. Each device supporting the DNP protocol must have a unique address within the collection of devices sending and receiving DNP messages.

22.1.6.1. Address Learning for DNP

The RX1500 implements both local and remote address learning for DNP.

A local Device Address Table is populated with DNP Addresses learned for local and remote DNP devices. Each DNP address is associated with either a local serial port or a remote IP address.

When a message with an unknown DNP source address is received on a local serial port, the DNP source address and serial port number are entered into the Device Address Table. When a message with an unknown DNP source address is received from the IP network, on the IP interface that is configured as the DNP learning interface, the DNP source address and the IP address of the sender are entered into the Device Address Table.

When a message with an unknown DNP destination address is received on a local serial port, the message is sent in a UDP broadcast out the network interface configured as the DNP learning interface. When a message with an unknown DNP destination address is received from the IP network, it is sent to all local serial ports configured as DNP ports.

UDP transport is used during the DNP address learning phase.

All learned addresses will be kept in the Device Address Table, which is saved in non-volatile memory, which makes it unnecessary to repeat the DNP address learning process across a RX1500 reboot or accidental power loss.

An aging timer is maintained per DNP address in the table, and is reset whenever a DNP message is sent to or received for the specified address.

This learning facility makes it possible to configure the DNP3 protocol with a minimum number of parameters: a TCP/UDP port number, a learning network interface and an aging timer.

22.1.6.2. DNP Broadcast Messages

DNP addresses 65521 through 65535 are reserved as DNP3 broadcast addresses. The RX1500 supports DNP3 broadcast messages. DNP broadcast messages received on local serial ports are transmitted to all IP Addresses in the DNP Device Address Table (whether learned or statically configured). When a DNP broadcast message is received from the IP network, it is transmitted on all local serial ports configured as DNP ports.

22.2. Serial Protocol Configuration

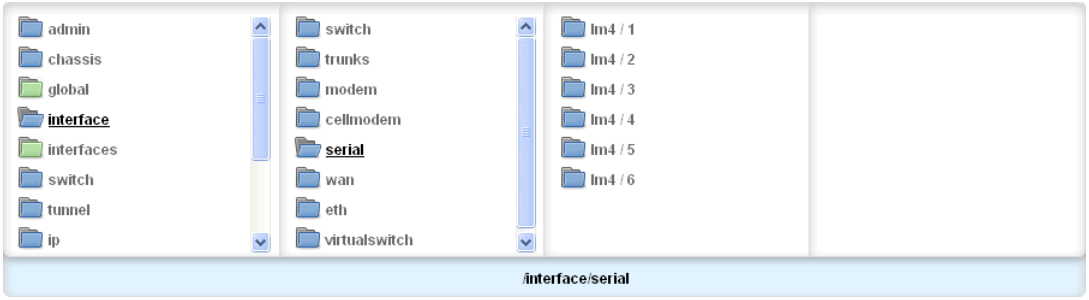

Figure 22.3. Serial Protocols menu

To display the Serial Protocols menu, navigate to *interface/serial*.

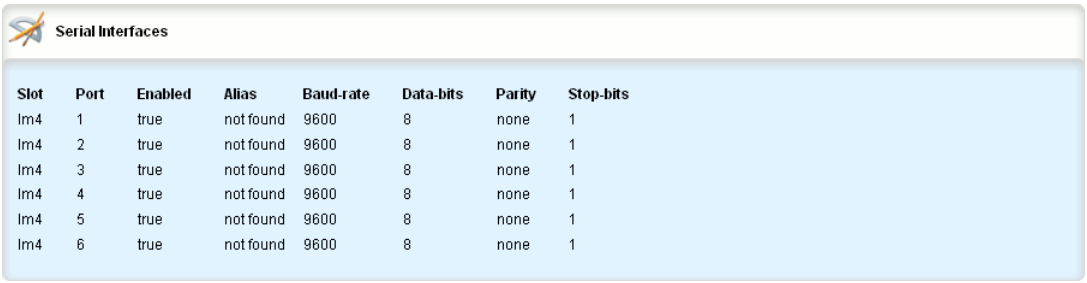
Serial Interfaces
Slot Port Enabled Alias Baud-rate Data-bits Parity Stop-bits
lm4 1 true not found 9600 8 none 1
lm4 2 true not found 9600 8 none 1
lm4 3 true not found 9600 8 none 1
lm4 4 true not found 9600 8 none 1
lm4 5 true not found 9600 8 none 1
lm4 6 true not found 9600 8 none 1

Figure 22.4. Serial Interfaces table

If data and ports have been configured, the Serial Interfaces table appears on the same screen as the *Serial Protocols* menu.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

enabled

Synopsis: *boolean*

Default: *true*

Provides the option to enable or disable this interface. When unchecked(i.e disabled), the interface will prevent all frames from being sent and received on that interface.

alias

Synopsis: *A string*

The SNMP alias name of the interface

22.2.1. Assigning Protocols

Select the type of protocol to assign to a serial port.

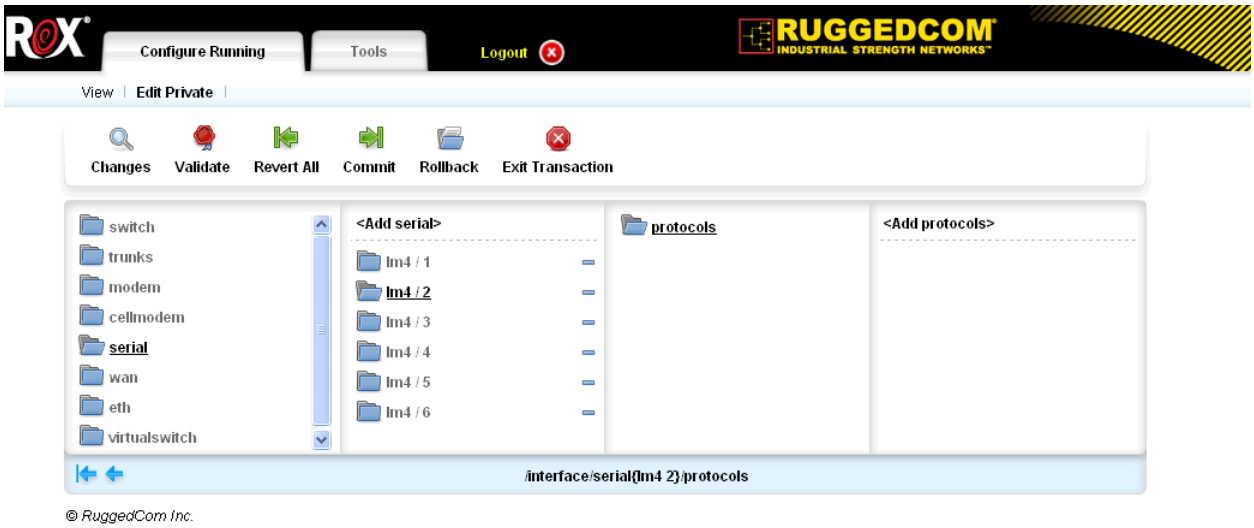

Figure 22.5. Adding a Protocol in the Edit Private screen

In Edit Private view, the <Add protocols> option can be clicked, which adds a protocol to a port.

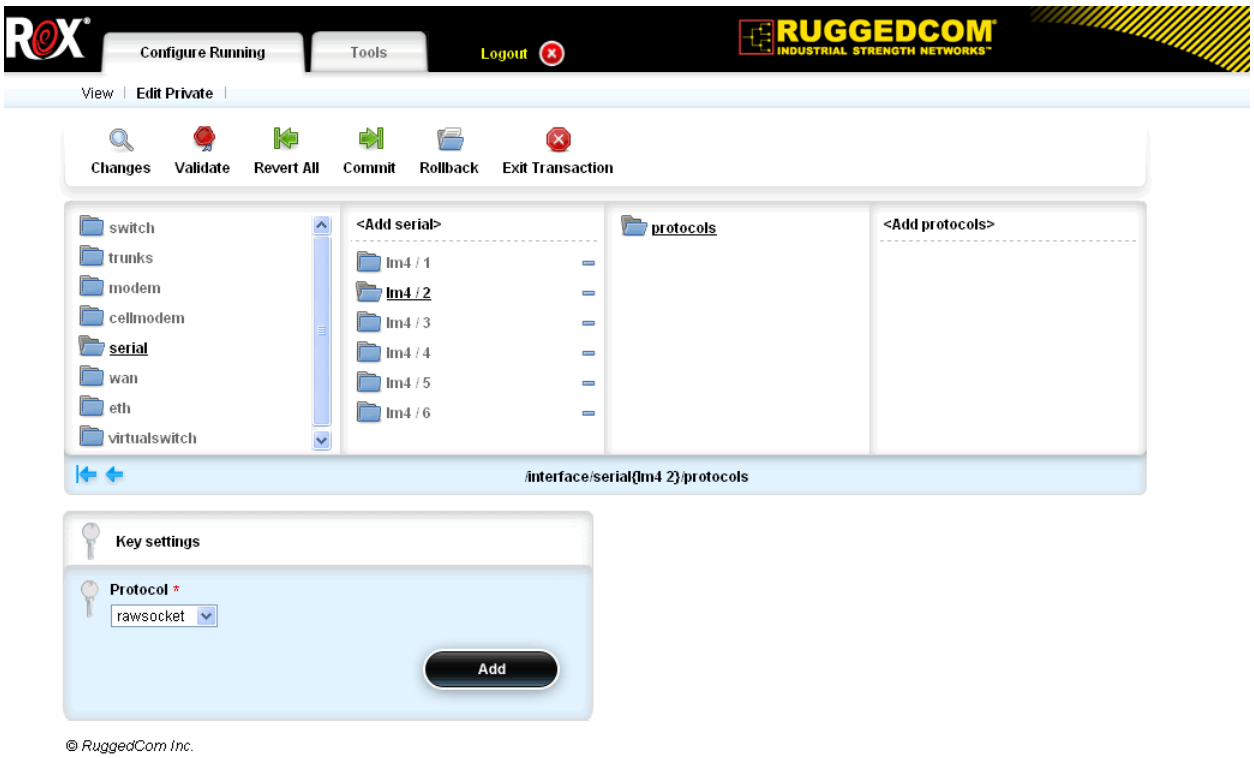

Figure 22.6. Selecting a Protocol Type in the Edit Private screen

Selecting a protocol type from the *Protocol* field in the Key Settings form associates a protocol with a serial port. Rawsocket, TcpModbus or DNP protocols can be selected. Unused ports should be left associated with "none". Changing an association immediately closes the calls of any protocols previously selected on the same port.

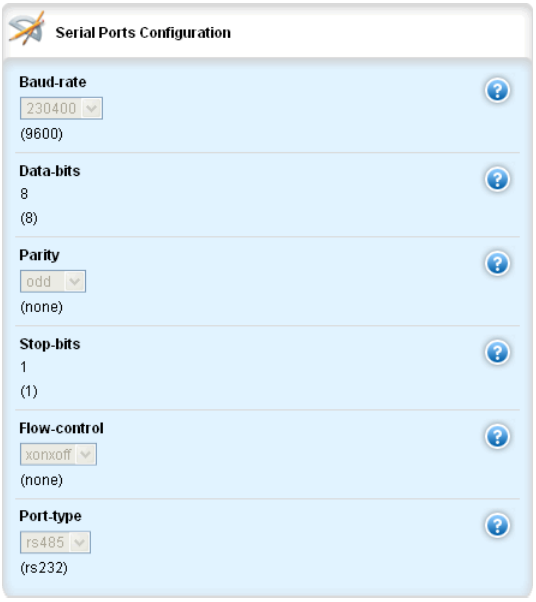

Figure 22.7. Serial Ports Configuration form

The Serial Interfaces form configures the serial settings and electrical protocol associated with a serial port. Changes are made immediately. To display this form, navigate to *interface/serial/{line module}*.

baud-rate

Synopsis: *string - one of the following keywords { 230400, 115200, 57600, 38400, 19200, 9600, 2400, 1200 }*

Default: *9600*

The baudrate selection of serial port

data-bits

Synopsis: *integer*

Default: *8*

The number of data bits

parity

Synopsis: *string - one of the following keywords { odd, even, none }*

Default: *none*

The parity of the serial port

stop-bits

Synopsis: *integer*

Default: *1*

The number of stop bits of the serial port

flow-control

Synopsis: *string - one of the following keywords { xonxoff, none }*

Default: *none*

Flow control of the serial port

port-type

Synopsis: *string - one of the following keywords { rs485, rs422, rs232 }*

Default: *rs232*

The type of serial port

Figure 22.8. Serial Protocols table

The Serial Protocols table displays the protocols configured. To display the Serial Interfaces table, navigate to *interface/serial/{line module}/protocols*.

protocol

Synopsis: *string - one of the following keywords { dnp, tcpmodbus, rawsocket }*

22.2.2. Setting Rawsockets

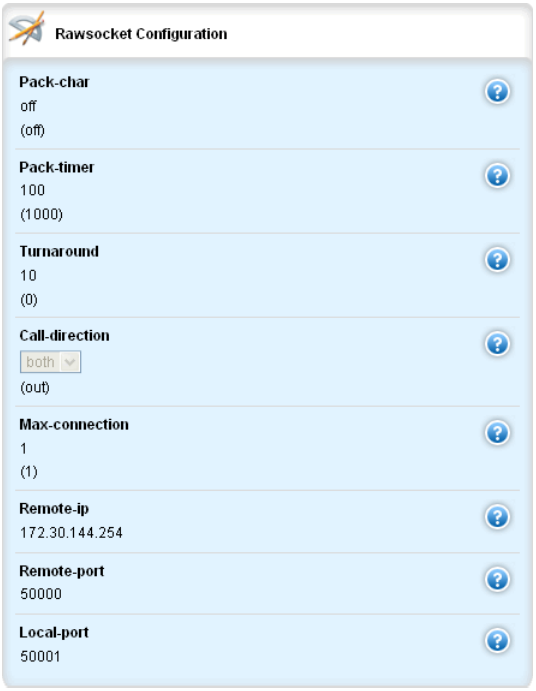

Figure 22.9. Rawsocket Configuration form

The Rawsocket Configuration form is used to configure the Raw Socket settings for each port. Changes are made immediately. To display the Rawsocket Configuration form, navigate to *interface/serial/{line module}/protocols/rawsocket/setrawsocket*.

pack-char

Synopsis: *string - the keyword { off }*

Synopsis: *integer*

Default: *off*

The numeric value of the ASCII character which will force forwarding of accumulated data to the network.

pack-timer

Synopsis: *integer*

Default: *1000*

The delay from the last received character until when data is forwarded

turnaround

Synopsis: *integer*

Default:

The amount of delay (if any) to insert between the transmissions of individual messages out the serial port

call-direction

Synopsis: *string - one of the following keywords { both, out, in }*

Default: *out*

Whether to accept an incoming connection, place an outgoing connection or do both

max-connection

Synopsis: *integer*

Default: *1*

The maximum number of incoming connections to permit when the call direction is incoming.

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address used when placing an outgoing connection.

remote-port

Synopsis: *integer*

The TCP destination port used in outgoing connections

local-port

Synopsis: *integer*

The local TCP port to use to accept incoming connections.

22.2.3. Setting TcpModbus

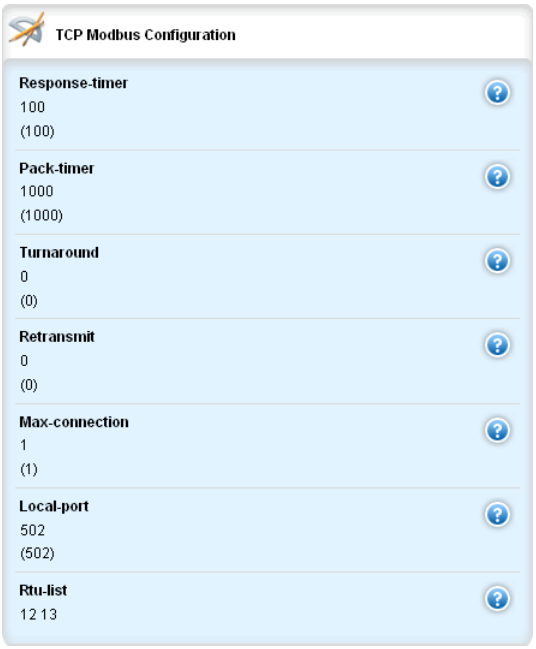

Figure 22.10. TCP Modbus Configuration form

The TCP Modbus Configuration form is used to configure the TcpModbus settings for each port. Changes are made immediately. To display the TCP Modbus Configuration form, navigate to *interface/serial/{line module}/protocols/tcpmodbus/settcpmodbus*.

response-timer

Synopsis: *integer*

Default: *100*

The maximum time from the last transmitted character of the outgoing poll until the first character of the response. If the RTU does not respond in this time, the poll will have been considered failed.

pack-timer

Synopsis: *integer*

Default: *1000*

The maximum allowable time to wait for a response to a Modbus

request to complete once it has started.

turnaround

Synopsis: *integer*

Default:

The amount of delay (if any) to insert after the transmissions of Modbus broadcast messages out the serial port.

retransmit

Synopsis: *integer*

Default:

The number of times to retransmit the request to the RTU before giving up.

max-connection

Synopsis: *integer*

Default: *1*

The maximum number of incoming connections.

local-port

Synopsis: *integer*

Default: *502*

The alternate local TCP port number. If this field is configured, a single connection (per serial port) may be made to this alternate port number. Note that TCP Modbus uses a default local port number of 502. There is no limit imposed on the number of connections to the default TCP port.

rtu-list

Synopsis: *string*

The ID of the RTU that is hooking up to the serial port.

The Modbus specification states the minimum time is about 640 character times at baud rates below 19200 Kbps and 256 char times + 192 ms at baud rates above 19200 Kbps. You may specify a larger value if you think your RTU will take longer to complete transmission than the calculated time.

22.2.4. Setting DNP

Figure 22.11. DNP Protocols Configuration form

The DNP Protocols Configuration form is used to configure the DNP settings for each port. To display the DNP Protocols Configuration form, navigate to *interface/serial/{line module}/protocols/dnp/setdnp*.

address-learning

Synopsis: *A string*

The interface to learn the RTU address from.

aging-timer

Synopsis: *integer*

Default: *1000*

The length of time which a learned DNP device in the Device Address Table may go without any DNP communication before it is removed from the table.

max-connection

Synopsis: *integer*

Default: *1*

The maximum number of incoming DNP connections.

Figure 22.12. DNP Device Address Table Configuration table

That Device Address table displays all currently known active DNP devices. To display the DNP Address Table Configuration form, navigate to *interface/serial/{line module}/protocols/dnp/setdnp/device-table*.

Figure 22.13. DNP Device Address Table Configuration form

To display the DNP Address Table Configuration form, navigate to *interface/serial/{line module}/protocols/dnp/setdnp/device-table/{number}*.

deviceAddress

Synopsis: *integer*

The local or remote DNP device address. The address may be that of a DNP device connected to a local serial port or one available via the serial port of a remote IP host.

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of the remote host that provides a connection to the DNP device with the configured address.Leave this field empty to forward DNP message that matches the configured address to local serial port

remote-device

Enable forwarding DNP message that matches the device address to remote-ip

22.3. Serial Protocol Statistics

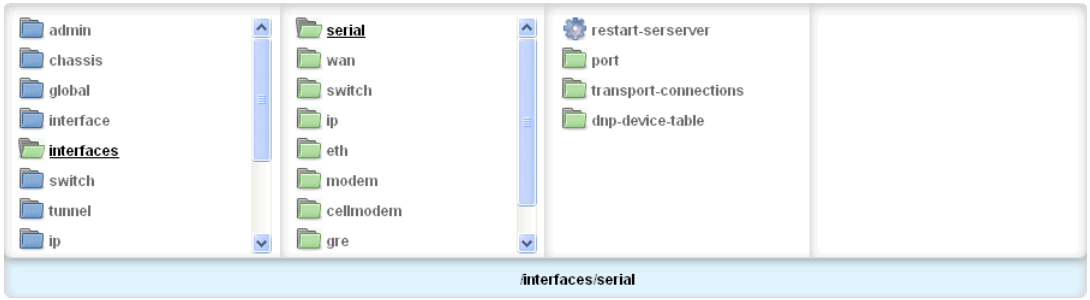

Figure 22.14. Serial Protocol Statistics menu

The Serial Protocol Statistics menu is accessible from the main menu under *interfaces/serial*.

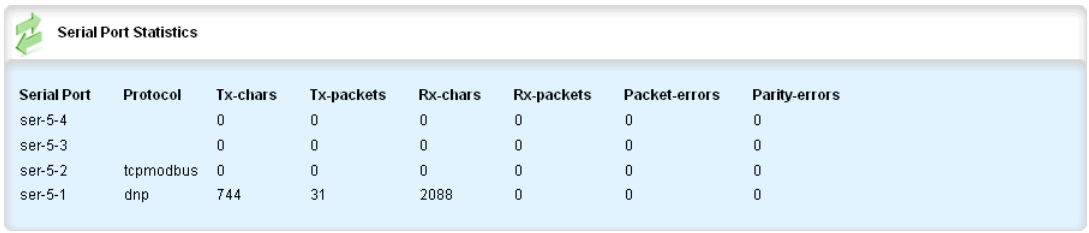

Serial Port	Protocol	Tx-chars	Tx-packets	Rx-chars	Rx-packets	Packet-errors	Parity-errors
ser-5-4		0	0	0	0	0	0
ser-5-3		0	0	0	0	0	0
ser-5-2	tcpmodbus	0	0	0	0	0	0
ser-5-1	dnp	744	31	2088	0	0	0

Figure 22.15. Serial Port Statistics table

To display the Serial Port Statistics table, navigate to *interfaces/serial/port*.

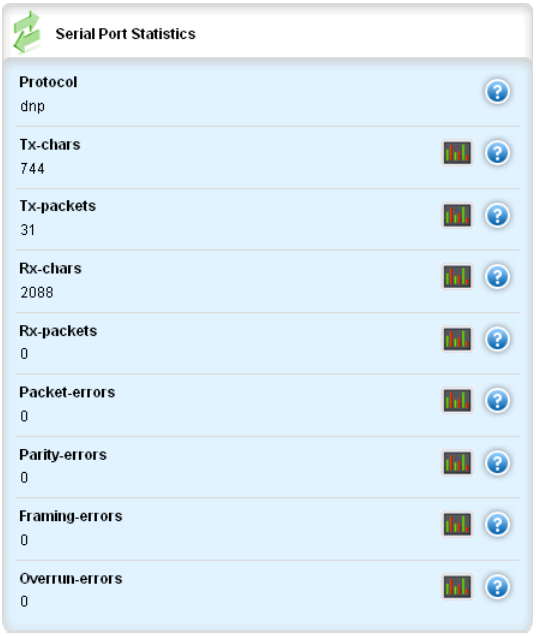

Figure 22.16. Serial Port Statistics form

To display the Serial Port Statistics form, navigate to *interfaces/serial/port* and then clicking on a linked submenu.

Serial Port

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

The serial interface name

media

Synopsis: *A string*

The type of port media { RS232 RS422 RS485 }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

speed

Synopsis: *string - one of the following keywords { 230.4K, 115.2K, 57.6K, 38.4K, 19.2K, 9.6K, 2.4K, 1.2K, 7.2M, 3.072M, 1.776M, 10G, 1G, 100M, 10M, 2.4M, 1.5M, auto }*

Speed (in Kilobits-per-second)

protocol

Synopsis: *A string*

The serial protocol assigned to this port

tx-chars

Synopsis: *unsigned integer*

The number of bytes transmitted over the serial port

tx-packets

Synopsis: *unsigned integer*

The number of packets transmitted over the serial port

rx-chars

Synopsis: *unsigned integer*

The number of bytes received by the serial port

rx-packets

Synopsis: *unsigned integer*

The number of packets received by the serial port

packet-errors

Synopsis: *unsigned integer*

The number of packet errors on this serial port

parity-errors

Synopsis: *unsigned integer*

The number of parity errors on this serial port

framing-errors

Synopsis: *unsigned integer*

The number of framing errors on this serial port

overrun-errors

Synopsis: *unsigned integer*

The number of overrun errors on this serial port

The Serial Port Statistics table and form present statistics of serial port activity and established connections. The Serial Port Statistics table provides a record for each active serial port. The number of historical received and transmitted characters as well as errors will be displayed.

22.3.1. Transport Connections

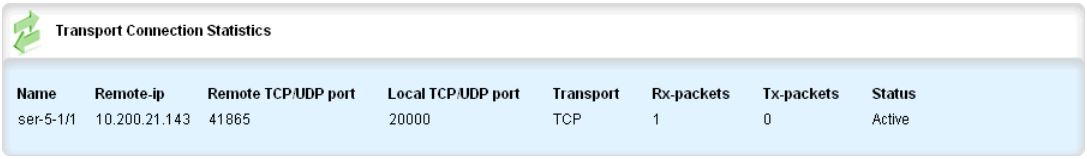

Figure 22.17. Transport Connections Statistics table

The Transport Connection Statistics table reflects established TCP connections. To display the Transport Connections Statistics table, navigate to *interfaces/serial/transport-connections*.

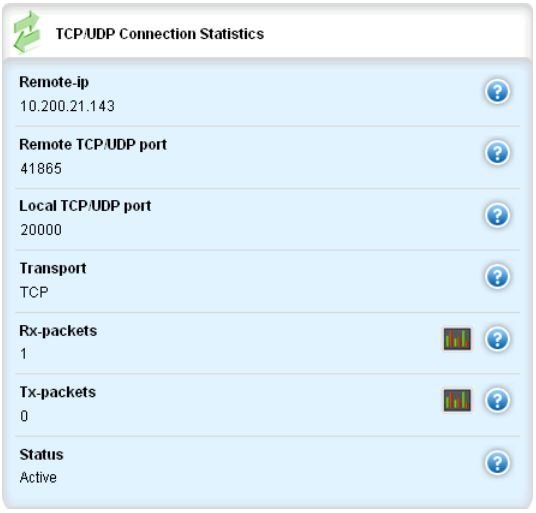

Figure 22.18. TCP/UDP Connection Statistics form

index

Synopsis: *A string*

The transport connection index

remote-ip

Synopsis: *A string*

The IP address of the remote serial server

Remote TCP/UDP port

Synopsis: *integer*

The port of the remote serial server

Local TCP/UDP port

Synopsis: *integer*

The local port for the incoming connection

transport

Synopsis: *A string*

The transport protocol (UDP or TCP) for this serial port

rx-packets

Synopsis: *unsigned integer*

The number of packets received from TCP/UDP

tx-packets

Synopsis: *unsigned integer*

The number of packets transmitted to TCP/UDP

target-port

Synopsis: *A string*

Target serial port

status

Synopsis: *A string*

The connection status of the serial port

22.4. Restarting the Serial Server

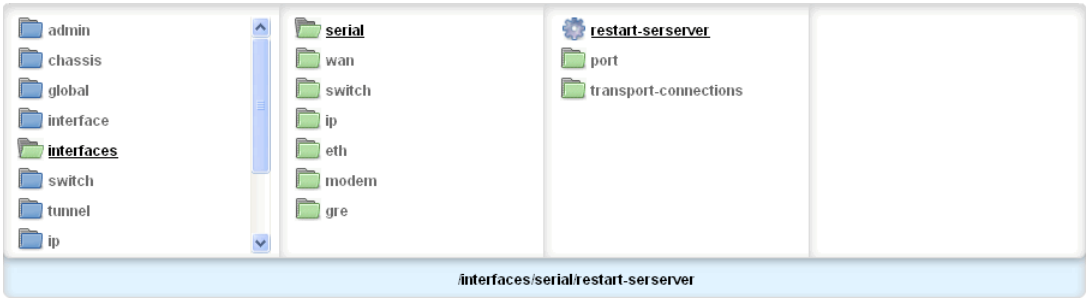

Figure 22.19. Restart Serserver menu

The path to the Restart Serserver menu is *interfaces/serial/restart-serserver*. To restart the serserver, click on the *restart-serserver* trigger action and the click the *Perform* button on the Trigger Action form.

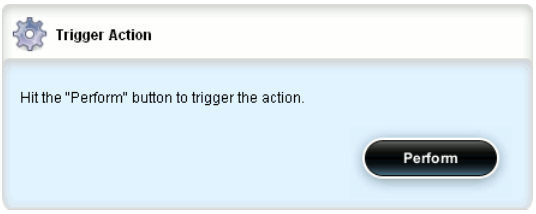

Figure 22.20. Restart Serserver Trigger Action

22.5. Resetting Ports

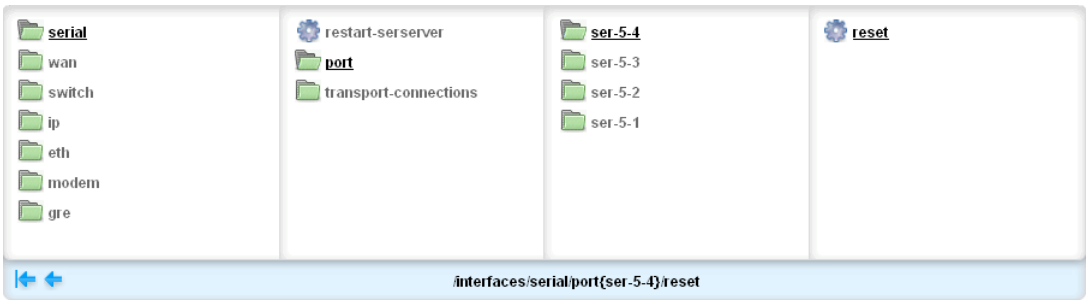

Figure 22.21. Reset Ports menu

The path to the Reset Ports menu is *interfaces/serial/port/{line module}/reset*. To reset the ports, click on the *reset* trigger action and then click the *Perform* button on the Reset Trigger Action form.

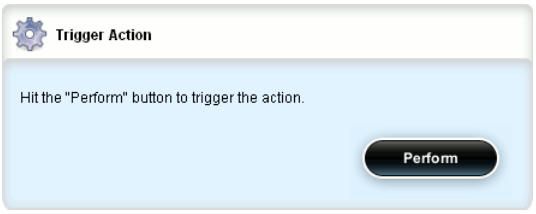

Figure 22.22. Reset Ports Trigger Action

23. WAN

23.1. T1/E1 Fundamentals

A T1 line is a communications circuit using the Digital Signal 1 (DS1) signalling scheme. DS1 allows 24 “timeslots” of 64 Kbps DS0 information, along with 8 Kbps of signalling information, to be multiplexed onto a 1544 Kbps circuit.

The 24 DS0s can be used individually as standalone channels, bonded into groups of channels, or can be bonded to form a single 1536 Kbps channel, referred to as a clear channel. Not all channels need be used. It is quite common to purchase a number of channels of 64Kbps bandwidth and leave the remainder unused; this is known as fractional T1.

The telephone network terminates the T1 line and maps each of the channels through the T1 network to a chosen T1 line. Individual and bonded DS0s from more than one remote T1 can be aggregated into a full T1 line. This is referred to as central site concentration.

The T1 line itself is referred to as the *physical interface*. Groups of DS0s form channels and the protocols that run on the channels are known as *logical interfaces*. The RuggedBackbone™ provides the ability to operate Frame Relay or PPP over logical interfaces.

An E1 line is a communications circuit conforming to European standards with 32 64 Kbps channels, one of which is usually reserved for signalling information.

23.1.1. Frame Relay

Frame Relay is a packet switching protocol for use over the WAN. The RuggedBackbone™ provides the ability to construct point-to-point IP network connections over Frame Relay.

Each Frame Relay interface provides a link between a local and a peer station. One of the stations must be configured as a Data Communications Equipment (DCE) device, often referred to as a Switch. The the peer station must be configured as a Data Terminal Equipment (DTE) device, often referred to as Customer Premises Equipment (CPE). The DCE is responsible for managing the link, advertising connections to the DTE, and switching packets between connections. The DTE raises individual connections and sends data on them. A frame relay switch is usually configured as DCE, and the RX1500 is configured as DTE.

When using a T1/E1 line to access a public Frame Relay provider, configure the router as a DTE.

Unlike PPP, a Frame Relay link can provide multiple connections. Each connection is identified by a Data Link Connection Identifier (DLCI) and must match at the DCE and DTE. The use of multiple connections can support meshed network interconnections and disaster recovery.

23.1.2. RX1500 and Frame Relay Encapsulation

The RX1500 supports IETF encapsulation for frame relay connections. When connecting the RX1500 to a Cisco router, you must configure the Cisco router to use IETF encapsulation. How you configure IETF encapsulation depends on the number of Data Link Connection Identifiers (DLCI) in use:

- When using a single physical Frame Relay interface and connecting to the RX1500 with one Data Link Connection Identifier (DLCI), enable IETF encapsulation as follows:

```
Cisco#conf terminal
            Cisco(config)#interface serial0/0
            Cisco(config-if)#encapsulation frame-relay IETF
            Cisco(config-if)#frame map ip ipaddress dlci broadcast
```

  Where ***ipaddress*** is the remote IP address, and ***dlci*** is the DLCI number.

- When using a single physical Frame Relay interface and connecting to the RX1500 with multiple DLCIs with mixed Cisco and IETF encapsulations, enable IETF encapsulation as follows. The text in [square brackets] indicates the type of encapsulation set by the command; do not type the text in [square brackets]:

```
Cisco(config)#interface serial0/0
            Cisco(config-if)#encap frame
            Cisco(config-if)#frame map ip ipaddress dlci{n} broadcast
 [Cisco]
            Cisco(config-if)#frame map ip ipaddress dlci{n} ietf
 broadcast [IETF]
              .
              .
              .
            Cisco(config-if)#frame map ip ipaddress dlci{n} ietf
 broadcast [IETF]
```

Where ***ipaddress*** is the remote IP address, and ***dlci{n}*** is a Data Link Connection Identifier number within the range of 16 to 1007.

23.2. WAN Configuration

Figure 23.1. WAN menu

To display the WAN menu, navigate to *interface/wan*. The WAN Slot Port Settings table appears on the same page as the WAN menu.

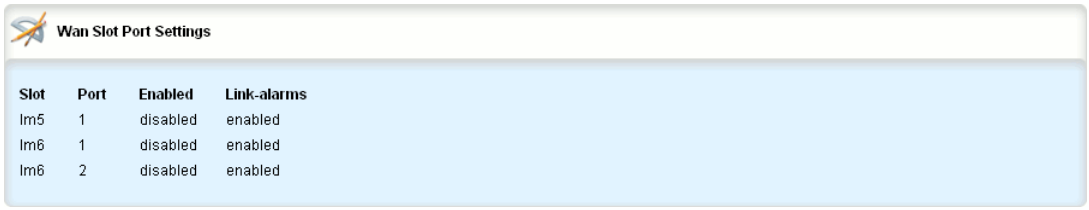
Wan Slot Port Settings
Slot Port Enabled Link-alarms
lm5 1 disabled enabled
lm6 1 disabled enabled
lm6 2 disabled enabled

Figure 23.2. Interface WAN Slot Port Settings table

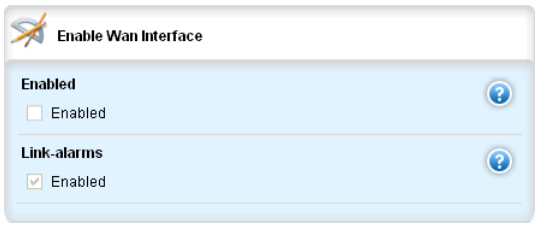

Figure 23.3. Enable WAN Interface form

The path to the Enable WAN Interface form is *interface/wan/{line module}*.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location for the WAN card.

port

Synopsis: *integer*

The port number on the WAN card.

enabled

Synopsis: *boolean*

Default: *false*

Enables this WAN port.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

alias

Synopsis: *A string*

The SNMP alias name of the interface

23.2.1. T1 Parameters

You can configure T1 parameters for a WAN port. The path to the T1 Parameters form is *interface/wan/{line module}/T1*.

Figure 23.4. T1 Parameters form

frame

Synopsis: *string - the keyword { esf }*

Default: *esf*

The frame format.

line-code

Synopsis: *string - the keyword { b8zs }*

Default: *b8zs*

The line encoding/decoding scheme.

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *normal*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

lbo

Synopsis: *string - one of the following keywords { 550-660ft, 440-550ft, 330-440ft, 220-330ft, 110-220ft, 0-110ft, 22.5db, 15db, 7.5db, 0db }*

Default: *0db*

Line Build Out: tunes the shape of the T1 pulses and adjusts their amplitude depending upon distances and the desired attenuation.

23.2.2. E1 Parameters

You can configure E1 Parameters for a WAN port. The path to the E1 Parameters form is *interface/wan/{line module}/E1*.

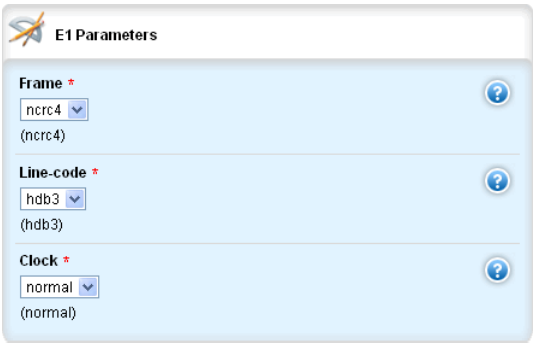

Figure 23.5. E1 Parameters form

frame

Synopsis: *string - one of the following keywords { crc4, ncrc4 }*

Default: *ncrc4*

The frame format.

line-code

Synopsis: *string - the keyword { hdb3 }*

Default: *hdb3*

A line encoding/decoding scheme.

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *normal*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

23.2.3. Configuring Protocols

The path to the T1 Channels and Associated Time Slots table is */interface/wan{line module and port}/t1/channel*.

Figure 23.6. T1 Channels and Associated Time Slots table

The path to the T1 Time Slots form is */interface/wan{line module and port}/t1/channel{channel id}*.

23.2.3.1. Adding Channels

You can configure one or more channels over one t1/e1 physical interface, and each channel can have one or more time slots.

> *A maximum of 24 timeslots (all of the timeslots) can be allocated to a channel. However, the same time slots cannot be assigned to two channels belonging to the same physical interface.*

To add a channel, follow these steps:

1. Enter edit mode and navigate to */interface/wan{lm6 2}/t1* or */interface/wan{lm6 2}/e1*.
2. Click the icon beside *t1* or *e1*.
3. Click channel and <Add channel>. The Key settings form appears.
4. In the Key settings form, enter a number in the range of 1 to 24 and click **Add**. The T1 Time Slots form appears.

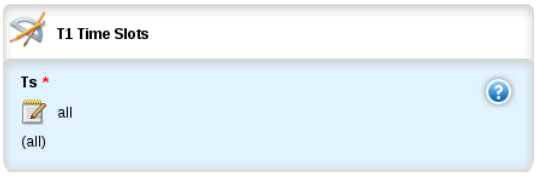

Figure 23.7. T1 Time Slots form

ts

Synopsis: *A string conforming to: "(all|[0-9,.]+)"*

Default: *all*

Time slots for this channel. Format: the string 'all', or a comma-separated list of numbers in the range of 1 to 24.

To specify a range of numbers, separate the start and end of the range with '..'

Example 1: 1,2,3 and 1..3 both represent time slots 1 through 3.

Example 2: 1,2,5..10,11 represents time slots 1, 2, 5, 6, 7, 8, 9, 10, and 11.

5. Commit the changes.

23.2.3.2. Adding Connections

After adding a channel, add connections to the channel. To add connections, enter edit mode and navigate to *interface/wan/{line module}/t1/{line module}/connection*.

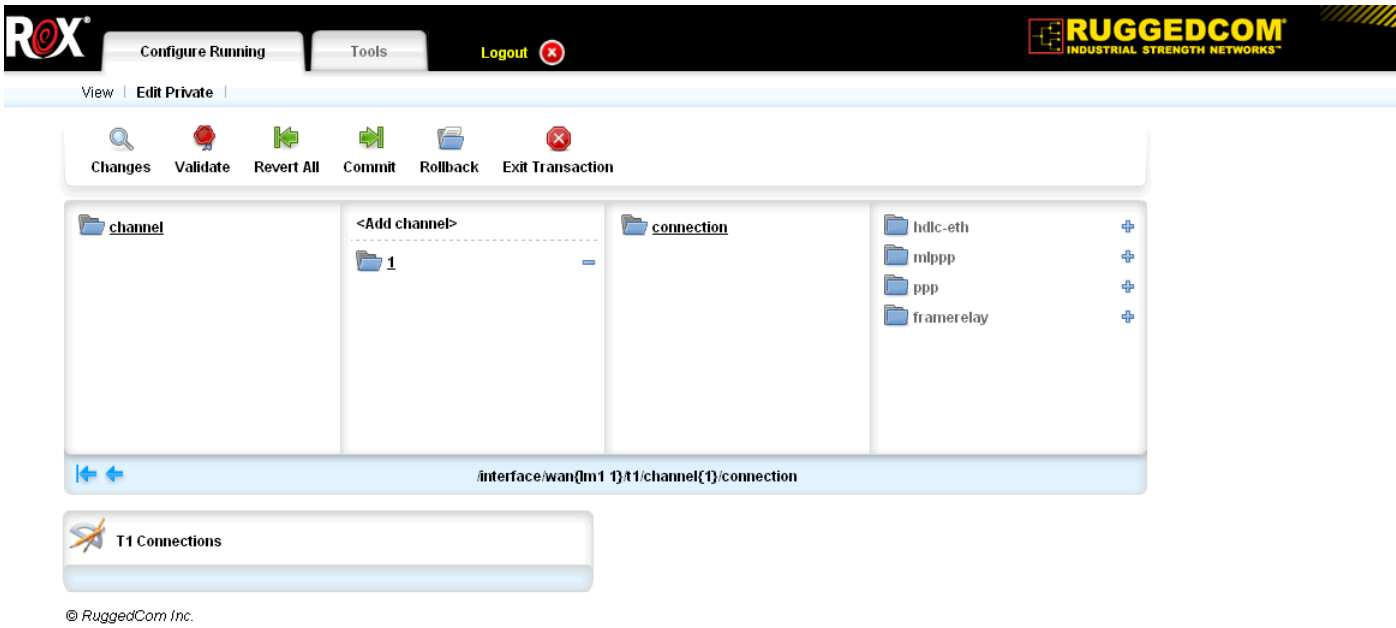

Figure 23.8. Adding a Connection

23.2.3.3. Configuring Frame Relay

From the *connection* submenu (see Figure 23.8, “Adding a Connection”), add a framerelay connection by clicking on the plus sign icon next to the *framerelay* submenu. Configure the parameters in the Frame Relay Parameter form.

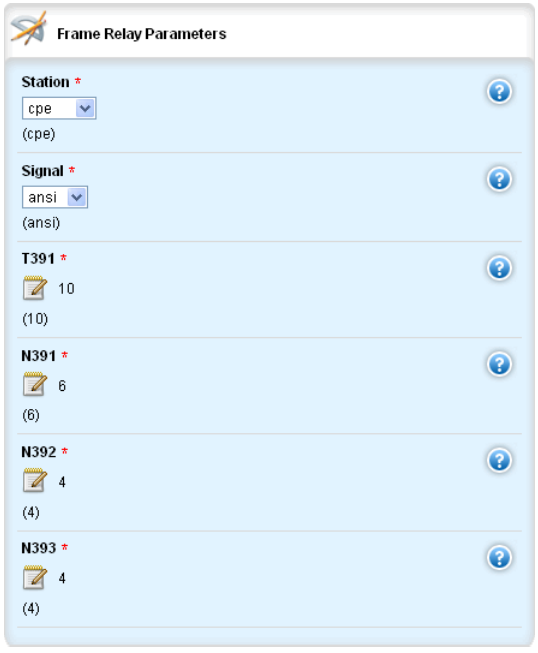

Figure 23.9. Frame Relay Parameter form

station

Synopsis: *string - one of the following keywords { switch, cpe }*

Default: *cpe*

The behavior of the frame relay connection, i.e. CPE (Customer Premises Equipment) or as a switch.

signal

Synopsis: *string - one of the following keywords { none, q933, lmi, ansi }*

Default: *ansi*

The frame relay link management protocol used.

t391

Synopsis: *integer*

Default: *10*

(Link Integrity Verification polling) Indicates the number of seconds between transmission of in-channel signaling messages. Valid for cpe.

n391

Synopsis: *integer*

Default: *6*

Defines the frequency of transmission of full status enquiry messages. Valid for CPE.

n392

Synopsis: *integer*

Default: *4*

The number of error events (enumerated by n393) for which the channel is declared inactive; valid for either cpe or Switch.

n393

Synopsis: *integer*

Default: *4*

The number of error events on the frame relay channel; valid for either

cpe or switch.

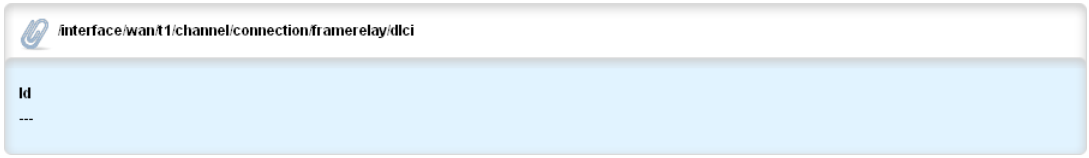

Figure 23.10. Connection Frame Relay DLCI table

This table displays the data link connection.

id

Synopsis: *integer*

DLCI (data link connection identifier) Number.

on-demand

This interface is up or down on demand of link fail over.

23.2.3.4. Configuring PPP

From the *connection* submenu (see Figure 23.8, “Adding a Connection”), add a PPP connection by clicking on the plus sign icon next to the *PPP* submenu. Click the *Commit* button and then click the *Exit Transaction* button. A PPP connection has now been added.

23.2.3.5. Configuring MLPPP

The PPP Multilink Protocol, also known as Multilink PPP or MLPPP, is defined in Internet RFC 1990. Its purpose is to combine two or more PPP links into a single “bundle” to provide more bandwidth for a point-to-point connection.

PPP Multilink must be supported on both sides of the link and may be used if there is more than one PPP link connecting the two endpoints. PPP works by multiplexing data on a per-packet basis to transmit across multiple PPP links. PPP uses sequence numbering to attempt to preserve the order of packets transmitted across the bundle.

ROX™ is capable of running PPP Multilink over two or more T1/E1 links, but is capable of defining only one MLPPP bundle.

For optimal PPP Multilink operation, ensure that each link in the MLPPP bundle has the same bandwidth: the number of time slots, the clocking mode, and the rate for each T1/E1 link used by PPP Multilink should be the same.

MLPPP interfaces use the following naming convention: **{physical interface}-mlppp-{bundle number}**. For example, **te1-mlppp-01** is an MLPPP interface on physical interface **te1** with bundle number **01**.

MLPPP interfaces use the following naming convention: **{physical interface}-mlppp-{bundle number}**. For example, **te1-mlppp-01** is an MLPPP interface on physical interface **te1** with bundle number **01**.

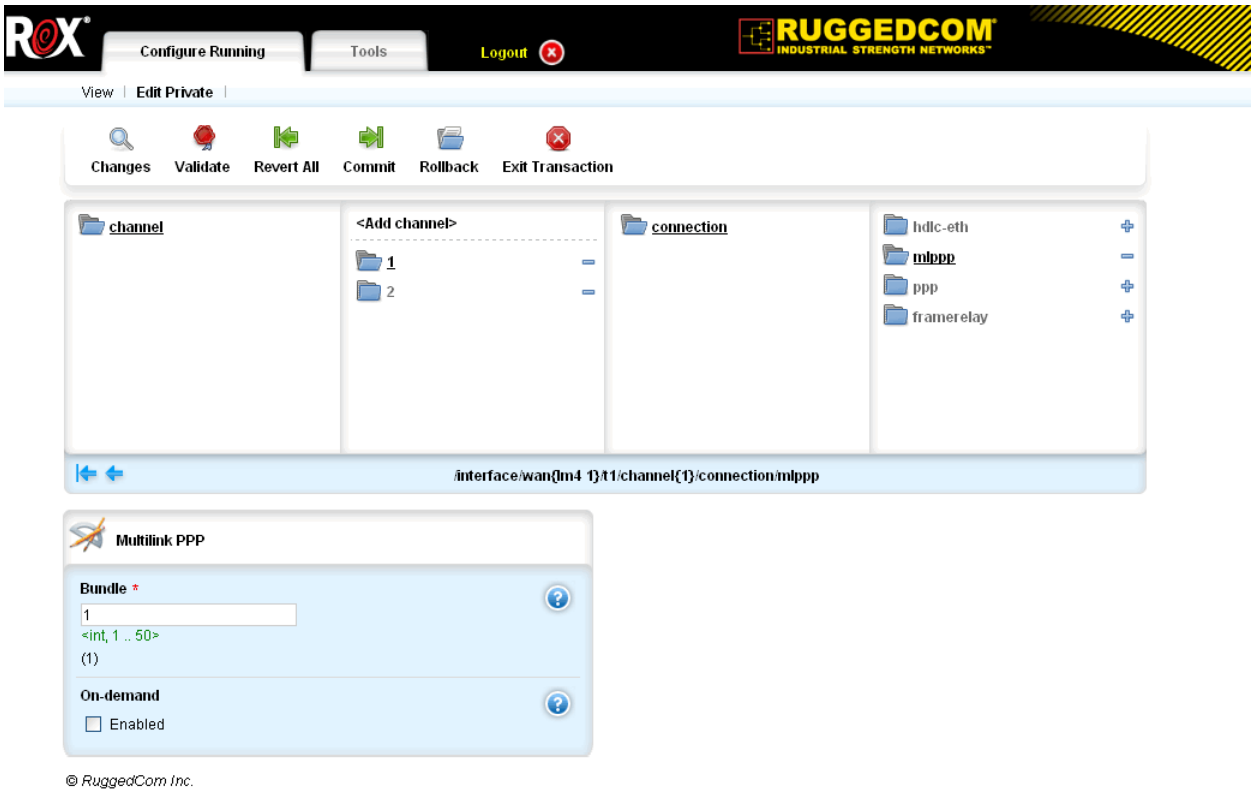

Figure 23.11. Adding an MLPPP Connection

From the *connection* submenu (see Figure 23.8, “Adding a Connection”), add an MLPPP connection by clicking the ✚ icon beside the *MLPPP* submenu. The Multilink PPP form appears.

Multilink PPP Form Parameters:

bundle

Synopsis: *integer*

Default: *1*

The bundle number

on-demand

This interface is up or down on demand of link fail over.

In the *Bundle* field, enter a bundle number (the default is 1). An MLPPP connection is added and an interface for this connection appears under the *ip* menu.

Figure 23.12. Adding IP and Remote Addresses

Navigate to *ip/{mlppp-interface}/ipv4* and click *Add address*. The Key settings form appears. In the *IPaddress* field, enter the IP address and click **Add**. The Addresses form appears.

In the *Peer* field, enter the remote IP address.

Click the **Commit** button and then click the **Exit Transaction** button.

You can add multiple PPP interfaces to a MLPPP link by configuring the same bundle number across all t1/e1 channels that are part of MLPPP.

23.2.3.6. Configuring HDLC-ETH

HDLC-ETH refers to Ethernet over an HDLC (High-Level Data Control) connection on a T1/E1 line. This connection passes Ethernet packets from a LAN through a T1/E1 line by creating a virtual switch containing one or more ethernet interfaces and an HDLC-ETH interface. For information on configuring a virtual switch, see Chapter 19, *Virtual Switch Bridging*.

A t1/e1 interface configured for HDLC-ETH works like a routable ethernet port (that is, fe-cm-1, switch.0001) which can be configured with an IP address and subnet mask. Because it acts like an ethernet port, you do not need to configure a peer IP address for an HDLC-ETH interface.

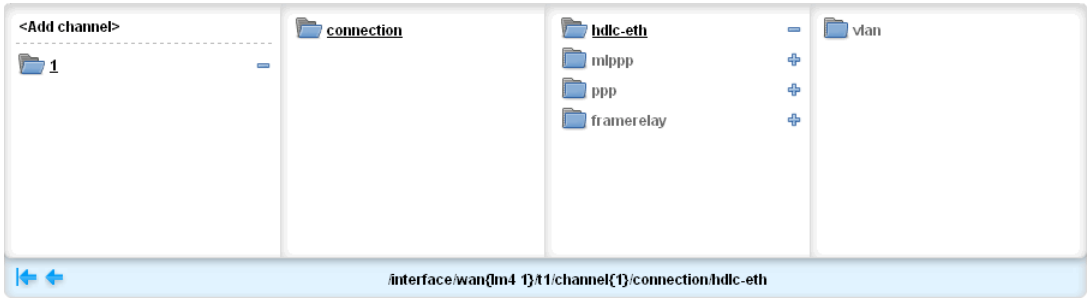

Figure 23.13. HDLC-ETH menu

Before adding an HDLC-ETH connection, you must first have a T1/E1 connection in place. For instructions on adding a T1/E1 connection, see Figure 23.8, “Adding a Connection”.

To add an HDLC-ETH connection, navigate to a T1/E1 connection at *interface/wan/{line module}/t1/channel{number}/connection/hdlc-eth* and click the icon beside the *hdlc-eth* submenu. An HDLC-ETH connection is added and the fields in the Ethernet Over HDLC Settings form become configurable.

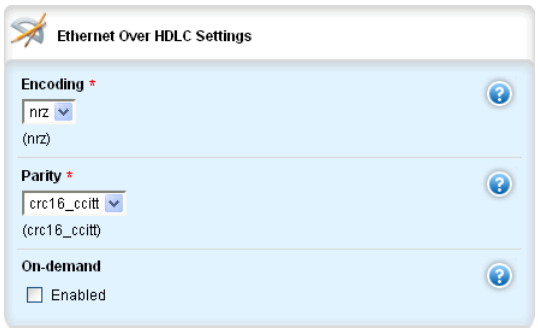

Figure 23.14. Ethernet Over HDLC Settings form

HDLC-ETH connections can be used with the default settings or can be configured in the Ethernet Over HDLC Settings form.

encoding

Synopsis: *string - the keyword { nrz }*

Default: *nrz*

HDLC encoding type

parity

Synopsis: *string - the keyword { crc16_ccitt }*

Default: *crc16_ccitt*

HDLC parity type

on-demand

This interface is up or down on demand of link fail over.

To add a VLAN, follow these steps:

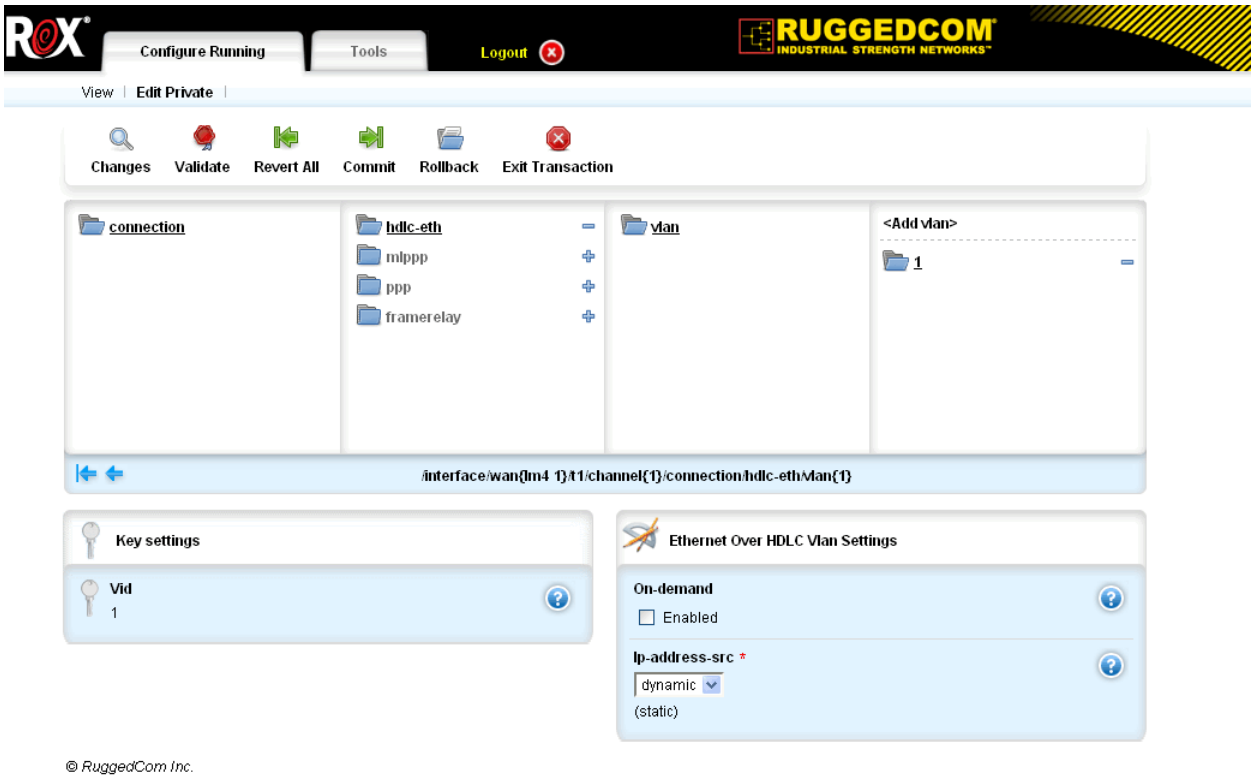

Figure 23.15. Adding a VLAN

- Click on the VLAN submenu and then on <Add vlan>. The Key settings form appears.
- On the Key settings form, enter a VLAN ID (VID) number and click **Add**. The Ethernet Over HDLC VLAN Settings form appears.
- Use the defaults or configure the settings in the Ethernet Over HDLC VLAN Settings form.

vid

Synopsis: *integer*

VLAN ID for this routable logical interface

on-demand

This interface is up or down on demand of link fail over.

ip-address-src

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP. The DYNAMIC option is a common case of a dynamically assigned IP address. It switches between BOOTP and DHCP until it gets the response from the relevant server. It must be static for non-management interfaces

23.2.3.7. Naming and IP Address Assignment for Logical Interfaces

All interface identifiers use the following naming convention:

te*N*-*[physical interface number]*-c*[channel identifier]*[*connection number*]

- **te*N*** identifies the interface as a WAN interface (for example, te1 represents t1/e1).
- ***[physical interface number]*** identifies the physical interface.

- **c*[channel number]*** identifies the channel number.
- ***[connection identifier]*** identifies the type of connection. The connection identifier can be any of the following:
  - **ppp**
  - **hdlc-eth**
  - **hdlc-eth** with VLAN ID
  - **mlppp** with a bundle number
  - ***fN*** for frame relay, where ***N*** is the Data Link Connection Identifier (DLCI), which is a four-digit number in the range of 0016 to 1007.

Examples:

- **te1-4-1c01** represents t1/e1 in slot 4 port 1, channel 1 is configured for hdlc-eth
- **te1-4-1c01.0012** represents t1/e1 in slot 4 port 1, channel 1 is configured for hdlc-eth with VLAN 12
- **te1-4-1c03ppp** represents t1/e1 in slot 4 port 1, channel 3 is configured for ppp
- **te1-4-1c04f0101** represents t1/e1 in slot 4 port 1, channel 4 is configured for FR DLCI 101
- **te1-mlppp-01** represents mlppp bundle 1

You can assign an IP Address for a t1/e1 logical interface in the IP submenu available from the main menu.

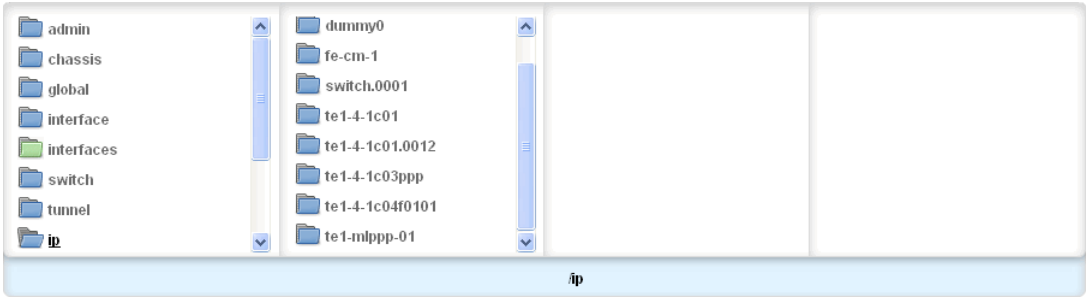

Figure 23.16. T1/E1 Interfaces under the IP submenu

> *Other than hdlc-eth, all other logical interfaces over t1/e1 are considered to be point to point links. Therefore, the only acceptable subnet mask for the point-to-point link is /32.*

23.2.4. Loopback Test

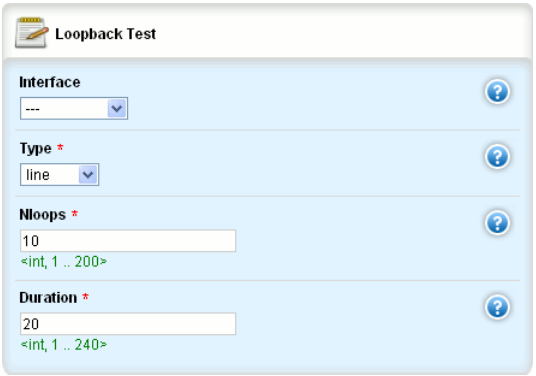

Figure 23.17. Loopback Test Forms

The path to the Loopback Test forms is *interfaces/wan/loopback*. In the Loopback Test form, select the *Interface* and *Type* and then set the *Nloops* and *Duration* parameters. To start a loopback test, click the **Perform** button on the trigger action form.

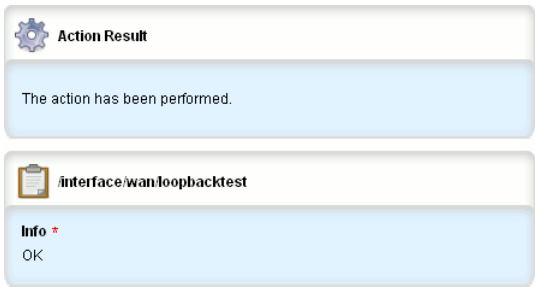

Figure 23.18. Loopbacktest Results

After launching the Loopback test, the Action Result form and the Loopbacktest form appear to confirm that the test has been performed and whether it was successful or not.

23.3. Statistics

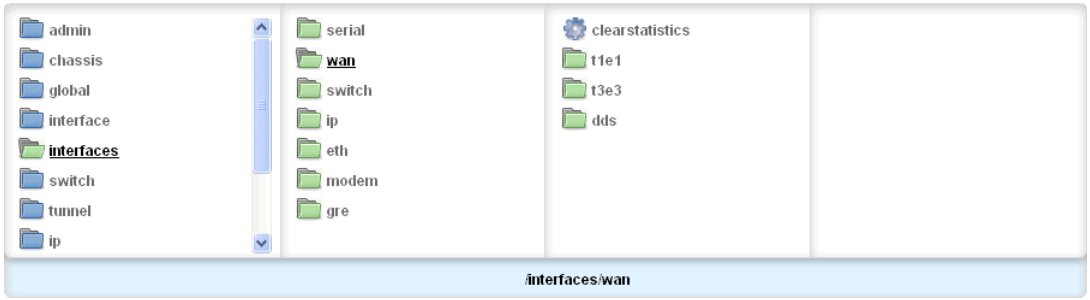

Figure 23.19. WAN Statistics menu

The WAN statistics are accessible via the *WAN Statistics* menu. The path to this menu is *interfaces/wan*.

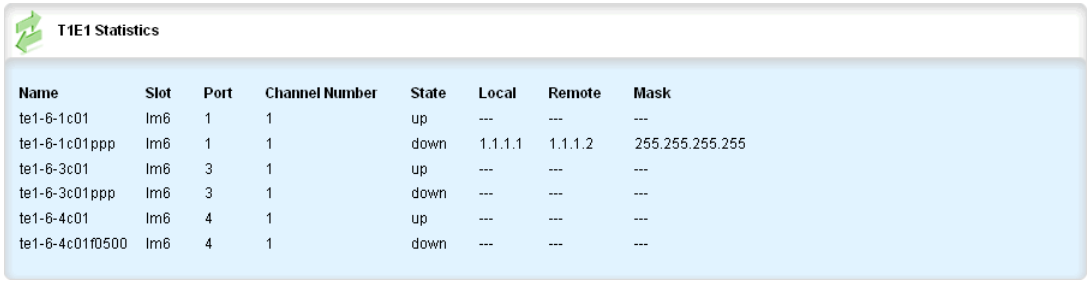

Name	Slot	Port	Channel Number	State	Local	Remote	Mask
te1-6-1c01	lm6	1	1	up	---	---	---
te1-6-1c01ppp	lm6	1	1	down	1.1.1.1	1.1.1.2	255.255.255.255
te1-6-3c01	lm6	3	1	up	---	---	---
te1-6-3c01ppp	lm6	3	1	down	---	---	---
te1-6-4c01	lm6	4	1	up	---	---	---
te1-6-4c01f0500	lm6	4	1	down	---	---	---

Figure 23.20. T1E1 Statistics table

The path to this table is *interfaces/wan/t1e1/{line module}/t1e1*.

23.3.1. Physical Layer-related Statistics

Figure 23.21. Receiving Errors Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/receive-error*.

Over Run

Synopsis: *unsigned integer*

The number of receiver overrun errors.

CRC Error

Synopsis: *unsigned integer*

The number of receiver CRC errors.

Abort

Synopsis: *unsigned integer*

The number of receiver abort errors.

Corruption

Synopsis: *unsigned integer*

The number of receiver corruption errors.

PCI Error

Synopsis: *unsigned integer*

The number of receiver PCI errors.

DMA Error

Synopsis: *unsigned integer*

The number of receiver DMA descriptor errors.

Length Error

Synopsis: *unsigned integer*

Length errors.

Frame Error

Synopsis: *unsigned integer*

Frame errors.

Figure 23.22. T1E1 Receiving Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/receive-stats*.

Frames

Synopsis: *unsigned integer*

The number of frames received.

Bytes

Synopsis: *unsigned integer*

The number of bytes received.

Frames Discarded as Link Inactive

Synopsis: *unsigned integer*

Received frames that were discarded (link inactive).

Figure 23.23. T1E1 Receiving Statistics Form 2

The path to this form is *interfaces/wan/t1e1/{line module}/receive*.

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of packets dropped.

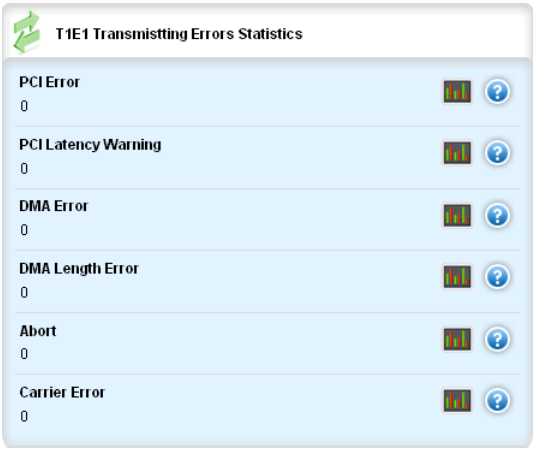

Figure 23.24. T1E1 Transmitting Errors Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/transmit-error*.

PCI Error

Synopsis: *unsigned integer*

The number of transmitter PCI errors.

PCI Latency Warning

Synopsis: *unsigned integer*

The number of transmitter PCI latency warnings.

DMA Error

Synopsis: *unsigned integer*

The number of transmitter DMA descriptor errors.

DMA Length Error

Synopsis: *unsigned integer*

The number of transmitter DMA descriptor length errors.

Abort

Synopsis: *unsigned integer*

Abort errors.

Carrier Error

Synopsis: *unsigned integer*

Carrier errors.

Figure 23.25. T1E1 Transmitting Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/transmit-stats*.

Frames

Synopsis: *unsigned integer*

The number of frames transmitted.

Bytes

Synopsis: *unsigned integer*

The number of bytes transmitted.

Realigned

Synopsis: *unsigned integer*

Transmits frames that were realigned.

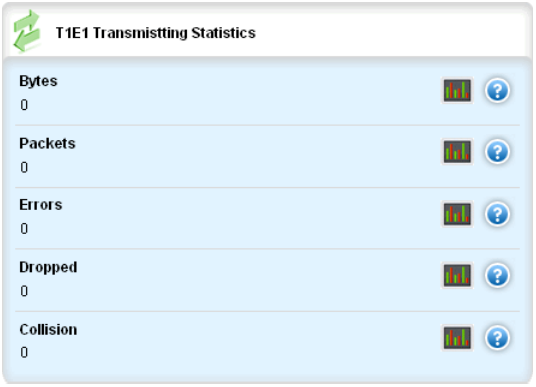

Figure 23.26. T1E1 Transmitting Statistics Form 2

The path to this form is *interfaces/wan/t1e1/{line module}/transmit*.

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets.

Dropped

Synopsis: *unsigned integer*

Number of packets dropped.

Collision

Synopsis: *unsigned integer*

Number of collisions detected.

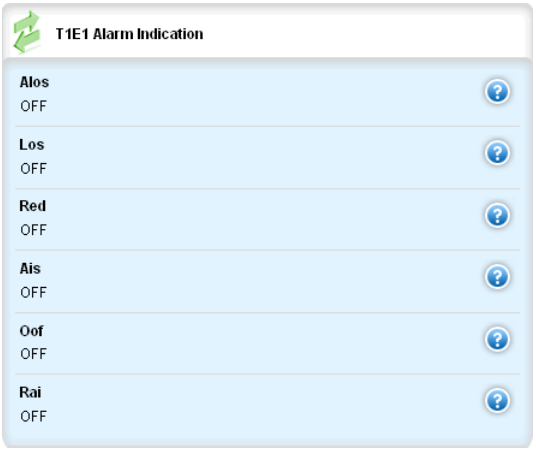

Figure 23.27. T1E1 Alarm Indication form

Alarm physical statistics are displayed in the T1E1 Alarm Indication form. The path to this form is *interfaces/wan/t1e1/{line module}/alarm*.

alos

Synopsis: *string*

ALOS (Loss of Signal) alarm.

los

Synopsis: *string*

LOS (Loss Of Signal) alarm.

red

Synopsis: *string*

RED (red alarm is a combination of a LOS or an OOF failure) alarm.

ais

Synopsis: *string*

AIS (Alarm Indication Signal) alarm.

oof

Synopsis: *string*

OOF (Out Of Frame) alarm.

rai

Synopsis: *string*

RAI (Remote Alarm Indication) alarm.

23.3.2. Protocol-related Statistics

23.3.2.1. PPP Statistics Summary

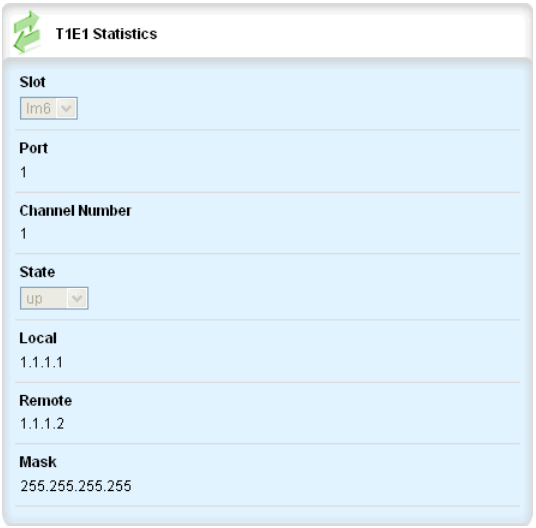

Figure 23.28. T1E1 Statistics form

The T1E1 Statistics form displays PPP statistics and physical statistics. The path to this form is *interfaces/wan/t1e1/{line module}*.

Figure 23.29. PPP Receiving Protocol Statistics form

The PPP Receiving Protocol Statistics form displays PPP receiving statistics. The path to this form is *interfaces/wan/t1e1/{line module}/ppp-stats*.

LCP

Synopsis: *unsigned integer*

The number of LCP (Link Control Protocol) packets.

PAP

Synopsis: *unsigned integer*

The number of PAP (Password Authentication Protocol) packets.

CHAP

Synopsis: *unsigned integer*

The number of CHAP (Challenge Handshake Authentication Protocol) packets.

IPCP

Synopsis: *unsigned integer*

The number of IPCP (Internet Protocol Control Protocol) packets.

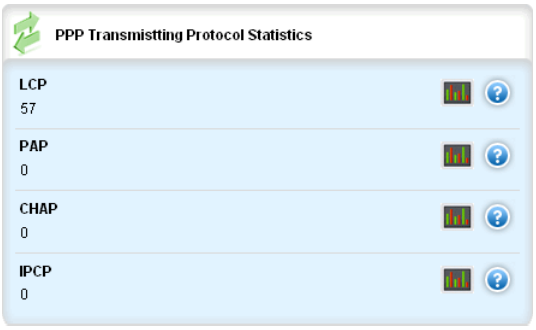

Figure 23.30. PPP Transmitting Protocol Statistics form

The PPP Receiving Protocol Statistics form displays PPP transmitting statistics. The path to this form is *interfaces/wan/t1e1/{line module}/ppp-stats*. Additional statistics forms can be found under this *ppp-stats* submenu.

LCP

Synopsis: *unsigned integer*

The number of LCP (Link Control Protocol) packets.

PAP

Synopsis: *unsigned integer*

The number of PAP (Password Authentication Protocol) packets.

CHAP

Synopsis: *unsigned integer*

The number of CHAP (Challenge Handshake Authentication Protocol) packets.

IPCP

Synopsis: *unsigned integer*

The number of IPCP (Internet Protocol Control Protocol) packets.

23.3.2.2. MLPPP Statistics

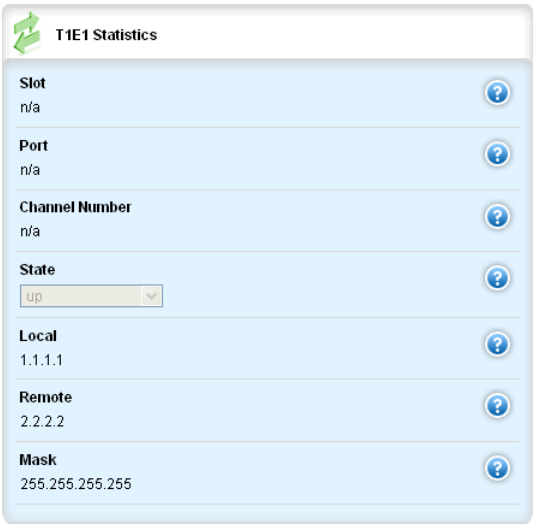

Figure 23.31. T1E1 Statistics form

The path to this form is *interfaces/wan/t1e1/{mlppp line module}*.

name

Synopsis: *A string*

Interface name.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string*

Line module name of the slot.

Port

Synopsis: *integer*

Synopsis: *string*

Port number on the slot.

Channel Number

Synopsis: *integer*

Synopsis: *string*

Channel number on the port.

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

Status of the interface.

local

Synopsis: *string*

Loacal IP address of the interface.

remote

Synopsis: *string*

Peer IP address.

mask

Synopsis: *string*

Netmask.

23.3.2.3. Frame-relay Statistics

Figure 23.32. Frame Relay Errors Packets Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/fr-stats/fr-error*.

Frame Length

Synopsis: *unsigned integer*

I-frames not transmitted after a tx. interrupt due to exessive frame length.

Throughput

Synopsis: *unsigned integer*

I-frames not transmitted after a tx. interrupt due to excessive throughput.

Length

Synopsis: *unsigned integer*

Received frames discarded as they were either too short or too long.

Unconfigured DLCI

Synopsis: *unsigned integer*

Discarded I-frames with unconfigured DLCI.

Format Error

Synopsis: *unsigned integer*

Discarded I-frames due to a format error.

No Response

Synopsis: *unsigned integer*

App. didn't respond to the triggered IRQ within the given timeout period.

Signalling Format Error

Synopsis: *unsigned integer*

Discarded In-channel Signalling frames due to a format error.

Invalid Send Sequence

Synopsis: *unsigned integer*

In-channel frames received with invalid Send Seq. Numbers received.

Invalid Receive Sequence

Synopsis: *unsigned integer*

In-channel frames received with invalid Receive Seq. Numbers received.

Unsolicited Response

Synopsis: *unsigned integer*

The number of unsolicited responses from the Access Node.

N391

Synopsis: *unsigned integer*

Timeouts on the T391 timer.

Consecutive T391

Synopsis: *unsigned integer*

Consecutive timeouts on the T391 timer.

N392 Error Threshold

Synopsis: *unsigned integer*

The number of times that the N392 error threshold was reached during N393 monitored events.

Figure 23.33. Frame Relay Controlling Packets Statistics form

The path to the Frame Relay Controlling Packets Statistics form and the Frame Relay Receiving Statistics form is *interfaces/wan/t1e1/{line module}/fr-stats*.

FSE

Synopsis: *unsigned integer*

Full Status Enquiry messages sent.

LIVSE

Synopsis: *unsigned integer*

Link Integrity Verification Status Enquiry messages sent.

FS

Synopsis: *unsigned integer*

Full Status messages received.

LIVS

Synopsis: *unsigned integer*

Link Integrity Verification Status messages received.

CPEI

Synopsis: *unsigned integer*

CPE initializations.

SSEQ

Synopsis: *unsigned integer*

The current Send Sequence Number.

RSEQ

Synopsis: *unsigned integer*

The current Receive Sequence Number.

N392

Synopsis: *unsigned integer*

The current N392 count.

N393

Synopsis: *unsigned integer*

The current N393 count.

Figure 23.34. Frame Relay Receiving Statistics form

Discard

Synopsis: *unsigned integer*

Received I-frames that were discarded due to inactive DLCI.

DEI

Synopsis: *unsigned integer*

I-frames received with the Discard Eligibility (DE) indicator set.

FECN

Synopsis: *unsigned integer*

I-frames received with the FECN bit set.

BECN

Synopsis: *unsigned integer*

I-frames received with the BECN bit set.

23.3.3. Clearing Statistics

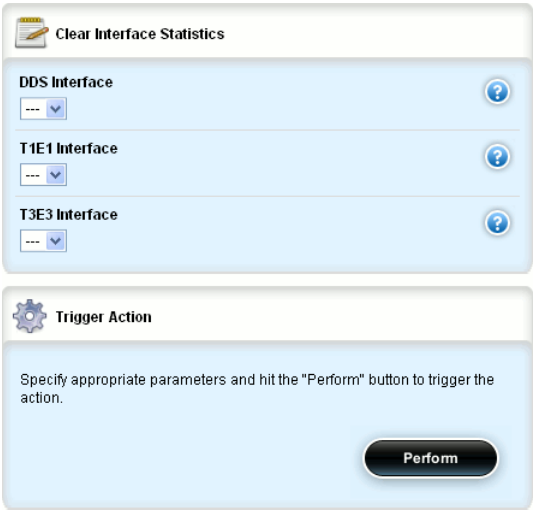

Figure 23.35. Clear Interface Statistics Form And Trigger Action

Statistics can be cleared by specifying the appropriate parameters in the Clear Interface Statistics form and then clicking the *Perform* button on the Trigger Action.

Figure 23.36. Clearstatistics Menu Action

The path to the Clear Statistics forms is *interfaces/wan/clearstatistics*.

23.4. DDS

Digital Data Services (DDS) is a North American digital transmission method. A DDS line operates synchronously at 56 Kbps or 64 Kbps over an unloaded 4-wire metallic-pair circuit.

A DDS line is typically a telephone-grade network connection and is often called the “local loop”. A Data Terminal Equipment (DTE) device is attached to the line and transmits data to the telephone company (TELCO), which routes the data to a remote DDS line. A short-haul synchronous-data line driver known as a CSU/DSU terminates the line and attaches to the DTE. The DSU part of the DSU/CSU manages

the format of the data signal. The CSU part of the DSU/CSU manages electrical levels and isolation, and provides loopback to the TELCO. A RuggedCom DDS port provides an integrated DTE, DSU, and CSU.

23.4.1. DDS Configuration

To configure DDS, you must first enable the WAN interface supporting DDS. See Section 23.4.1.1, “Enabling the DDS WAN Interface”.

Under */interface/wan{wan slot and port}/dds* you can configure the following:

- set the DDS parameters. See Section 23.4.1.2, “Setting DDS Parameters”.
- set the DDS PPP connection parameters. See Section 23.4.1.3, “Setting DDS PPP Connection Parameters”.
- set the DDS frame relay and DLCI parameters. See Section 23.4.1.4, “Setting DDS Frame Relay Parameters”.

Under interfaces/wan, you can also:

- view and clear DDS statistics. See Section 23.4.2, “Viewing and Clearing DDS Statistics”.
- perform a DDS loopback test. See Section 23.4.1.5, “Performing a DDS Loopback Test”.

23.4.1.1. Enabling the DDS WAN Interface

Before setting DDS parameters, you must first enable the WAN interface supporting DDS.

To enable the WAN DDS interface:

- In edit mode, navigate to */interface/wan{wan slot and port}*.
- On the *Enable Wan Interface* form, select **Enabled**.

Figure 23.37. Enable Wan Interface form

- Commit the changes.

23.4.1.2. Setting DDS Parameters

To set DDS parameters, enter edit mode and navigate to */interface/wan{wan slot and port}/dds/ddsparams*.

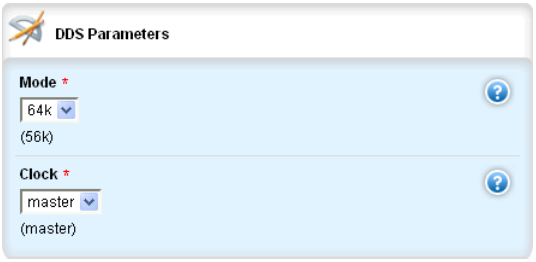

Figure 23.38. DDS Parameters form

mode

Synopsis: *string - one of the following keywords { 64k, 56k }*

Default: *56k*

DDS speed mode (kbps).

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *master*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

23.4.1.3. Setting DDS PPP Connection Parameters

To set DDS PPP connection parameters, enter edit mode and navigate to */interface/wan{wan slot and port}/dds/connection/ppp*.

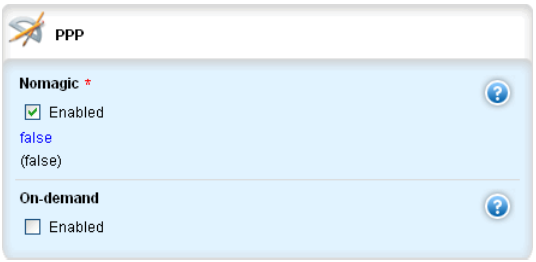

Figure 23.39. PPP form

nomagic

Synopsis: *boolean*

Default: *false*

Disables the Magic Number.

on-demand

This interface is up or down on demand of link fail over.

For more information on the On-demand function, see Section 41.3.4, “Link Backup On Demand”.

23.4.1.4. Setting DDS Frame Relay Parameters

You configure DDS frame relay parameters in two locations:

- */interface/wan{wan slot and port}/dds/connection/framerelay* configures general frame relay parameters

- */interface/wan{wan slot and port}/dds/connection/framerelay/dlci* configures the Data Link Connection Identifier (DLCI) parameters.

To set the DDS frame relay parameters:

- Enter edit mode and navigate to */interface/wan{wan slot and port}/dds/connection/framerelay*.

Figure 23.40. Frame Relay Parameters form

station

Synopsis: *string - one of the following keywords { switch, cpe }*

Default: *cpe*

The behavior of the frame relay connection, i.e. CPE (Customer Premises Equipment) or as a switch.

signal

Synopsis: *string - one of the following keywords { none, q933, lmi, ansi }*

Default: *ansi*

The frame relay link management protocol used.

t391

Synopsis: *integer*

Default: *10*

(Link Integrity Verification polling) Indicates the number of seconds between transmission of in-channel signaling messages. Valid for cpe.

t392

Synopsis: *integer*

Default: *16*

(Verification of polling cycle) Indicates the expected number of seconds between reception of in-channel signaling messages transmitted by cpe. Valid for Switch.

n391

Synopsis: *integer*

Default: *6*

Defines the frequency of transmission of full status enquiry messages. Valid for CPE.

n392

Synopsis: *integer*

Default: *4*

The number of error events (enumerated by n393) for which the channel is declared inactive; valid for either cpe or Switch.

n393

Synopsis: *integer*

Default: *4*

The number of error events on the frame relay channel; valid for either

cpe or switch.

eektype

Synopsis: *string - one of the following keywords { request, off }*

Default: *off*

Enables or disables EEK function.

eektimer

Synopsis: *integer*

Default: *5*

The number of seconds to wait before the next EEK message is sent.

To set the DDS frame relay DLCI parameters:

- Enter edit mode and navigate to */interface/wan{wan slot and port}/dds/connection/framerelay*.
- Beside the **framerelay** link, click the ✚ icon and click **<Add dlci>**.
- On the *Key settings* form, enter a number in the range of 16 to 1007 and click **Add**.
- On the *On demand* form, set the On-demand parameter. For more information on the On-demand function, see Section 41.3.4, “Link Backup On Demand”.

23.4.1.5. Performing a DDS Loopback Test

To perform a DDS loopback test, the remote equipment must be able to loop to allow verification of the entire line. If the remote equipment is an RX1000 or RX1500, starting a line loopback will verify both cards and the line. Note that DDS has no standard for performing digital loopback.

To perform a DDS loopback test:

- Navigate to */interfaces/wan/loopback*.
- On the *Loopback Test* form, set the test parameters.

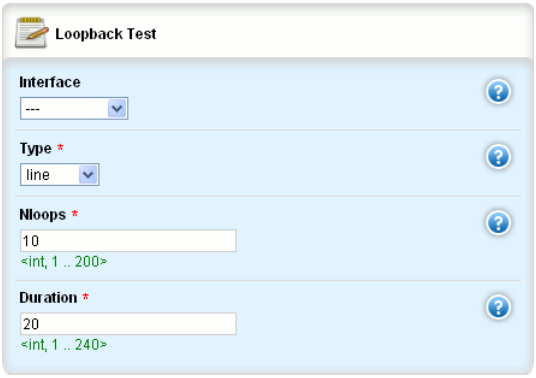

Figure 23.41. Loopback Test form

Interface

Physical interface name.

Type

The loopback type.

Nloops

The number of loops.

Duration

The number of seconds required to run the test.

- On the *Trigger Action* form, click **Perform**.

23.4.2. Viewing and Clearing DDS Statistics

DDS statistics are available when at least one logical interface is configured.

The main DDS statistics menu is available at */interfaces/wan/dds*.

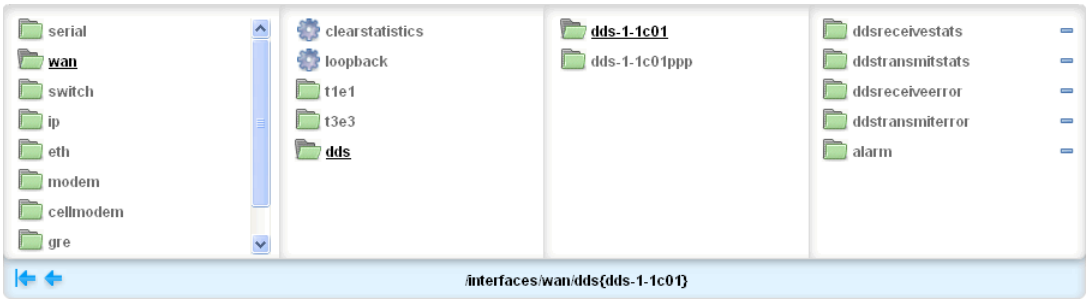

Figure 23.42. DDS Statistics menu

You can view DDS statistics by physical and logical connection. To view DDS physical connection statistics, navigate to:

Path	DDS Physical Connection Statistics
/interfaces/wan/dds{dds physical connection}/ddsreceivestats	Displays DDS physical connection receive statistics.
/interfaces/wan/dds{dds physical connection}/ddstransmitstats	Displays DDS physical connection transmit statistics.
/interfaces/wan/dds{dds physical connection}/ddsreceiveerror	Displays DDS physical connection receive error statistics.
/interfaces/wan/dds{dds physical connection}/ddstransmiterror	Displays DDS physical connection transmit error statistics.
/interfaces/wan/dds{dds physical connection}/alarm	Displays DDS physical connection alarm statistics.

Table 23.1. Paths to DDS Physical Connection Statistics Forms

To view DDS logical connection statistics, navigate to:

Path	DDS Logical Connection Statistics
/interfaces/wan/dds{dds logical connection}/receive	Displays DDS logical connection receive statistics.
/interfaces/wan/dds{dds logical connection}/transmit	Displays DDS logical connection transmit statistics.
/interfaces/wan/dds{dds logical connection}/protocolstats/ppp-stats	Displays DDS PPP protocol statistics.

Table 23.2. Paths to DDS Logical Connection Statistics Forms

23.4.2.1. Clearing DDS Statistics

To clear DDS statistics:

- Navigate to */interfaces/wan/clearstatistics*.
- On the *Clear Interface Statistics* form, set the clear statistics parameters.

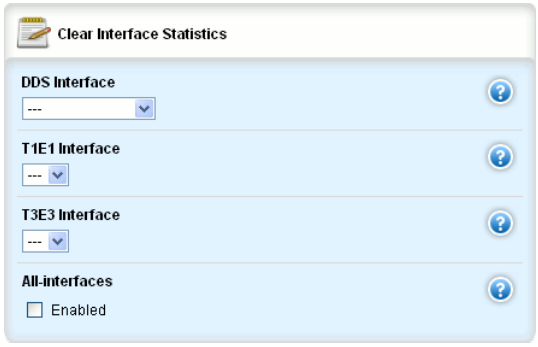

Figure 23.43. Clear Interface Statistics form

DDS Interface

Select the DDS interface for which to clear statistics.

T1E1 Interface

Select the T1E1 interface for which to clear statistics.

T3E3 Interface

Select T3E3 interface for which to clear statistics.

All-interfaces

Clear statistics for all WAN interfaces.

- On the *Trigger Action* form, click **Perform**.

24. Port Security

ROX™ Port Security provides the following features:

- Authorizing network access using Static MAC Address Table.
- Authorizing network access using IEEE 802.1X authentication.
- Configuring IEEE 802.1X authentication parameters.
- Detecting port security violation attempt and performing appropriate actions.

24.1. Port Security Operation

Port Security, or Port Access Control, provides the ability to filter or accept traffic from specific MAC addresses.

Port Security works by inspecting the source MAC addresses of received frames and validating them against the list of MAC addresses authorized on the port. Unauthorized frames will be filtered and, optionally, the port that receives the frame will be shut down permanently or for a period of time.

Frames to unknown destination addresses will not be flooded through secure ports.

Port security is applied at the edge of the network in order to restrict admission to specific devices. Do not apply port security on core switch connections.

ROX™ supports the MAC address authorization methods described below:

24.1.1. Static MAC address-based authorization

- With this method, the switch validates the source MAC addresses of received frames against the contents in the Static MAC Address Table.
- ROX™ also supports a highly flexible Port Security configuration which provides a convenient means for network administrators to use the feature in various network scenarios.
- A Static MAC address can be configured without a port number being explicitly specified. In this case, the configured MAC address will be automatically authorized on the port where it is detected. This allows devices to be connected to any secure port on the switch without requiring any reconfiguration.
- The switch can also be programmed to learn (and, thus, authorize) a preconfigured number of the first source MAC addresses encountered on a secure port. This enables the capture of the appropriate secure addresses when first configuring MAC address-based authorization on a port. Those MAC addresses are automatically inserted into the Static MAC Address Table and remain there until explicitly removed by the user.

24.1.2. IEEE 802.1X Authentication

The IEEE 802.1X standard defines a mechanism for port-based network access control and provides a means of authenticating and authorizing devices attached to LAN ports.

Although 802.1X is mostly used in wireless networks, this method is also implemented in wired switches.

The 802.1X standard defines three major components of the authentication method: Supplicant, Authenticator and Authentication server.

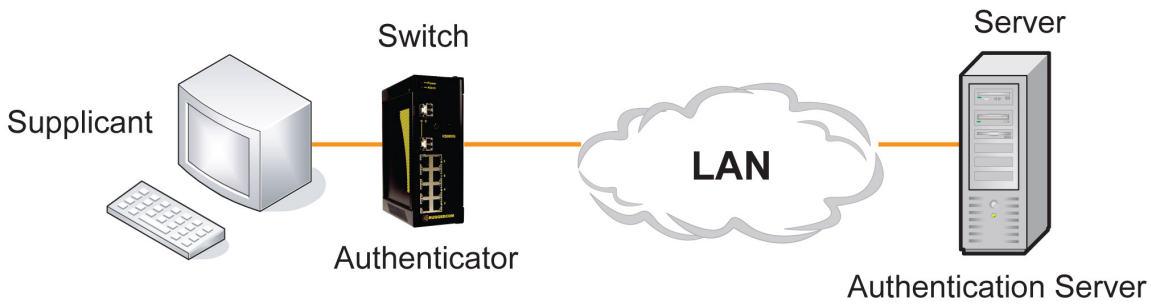

Figure 24.1. 802.1X General Topology

ROX™ supports the Authenticator component.

802.1X makes use of Extended Authentication Protocol (EAP) which is a generic PPP authentication protocol and supports various authentication methods. 802.1X defines a protocol for communication between the Supplicant and the Authenticator, EAP over LAN (EAPOL).

RuggedBackbone™ communicates with the Authentication Server using EAP over RADIUS.

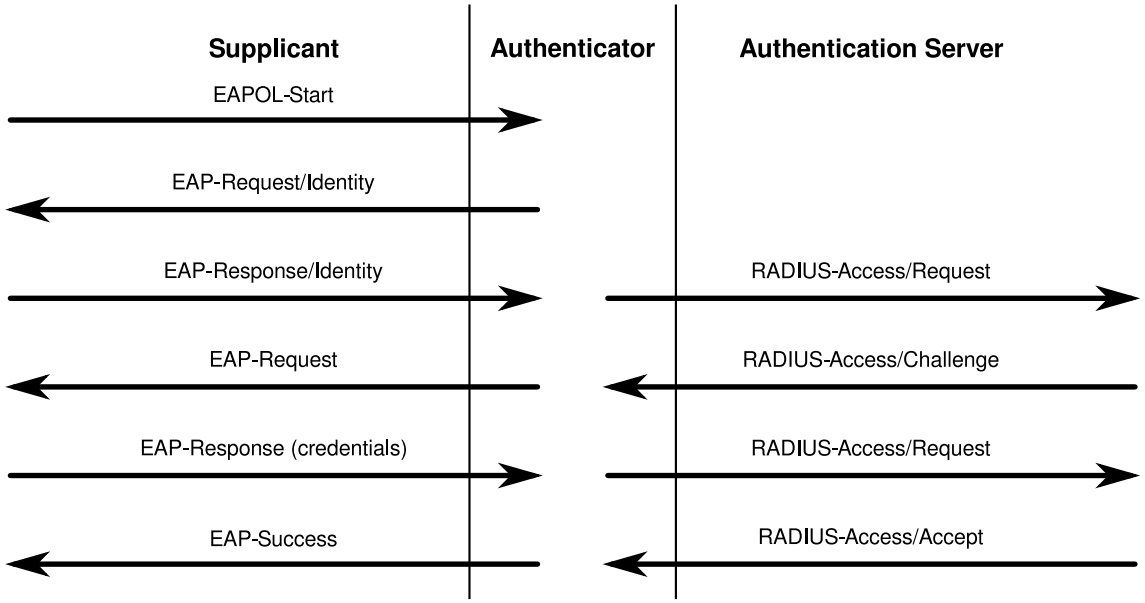

Figure 24.2. 802.1X Packet Exchange

The switch supports authentication of one host per port.

If the host's MAC address is configured in the Static MAC Address Table, it will be authorized, even if the host authentication is rejected by the authentication server.

24.1.2.1. RADIUS

Figure 24.3. Port Security RADIUS Primary form

The path to the Port Security RADIUS Primary form is *switch/port-security/radius*.

Figure 24.4. Port Security RADIUS Secondary form

The path to the Port Security RADIUS Secondary form is *switch/port-security/radius*.

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

UDP Port

Synopsis: *integer*

Default: *1812*

The IPv4 port of the server

password

Synopsis: *"AES CFB128"-encrypted string*

The password of the server

RADIUS servers can be configured for port security. For more information on RADIUS, please see Section 10.1, “RADIUS”

24.2. Port Security Configuration

To display the Port Security menu, Port Security form, and 802.1x Parameters form, navigate to *interface/switch/{line module}/port-security*.

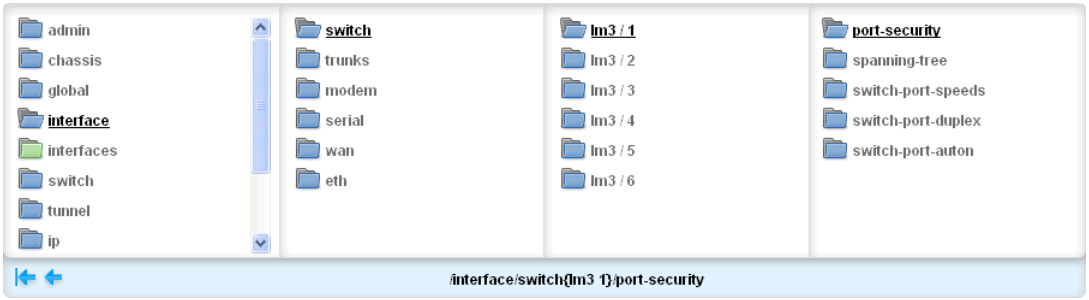

Figure 24.5. Port Security menu

24.2.1. Port Security Parameters

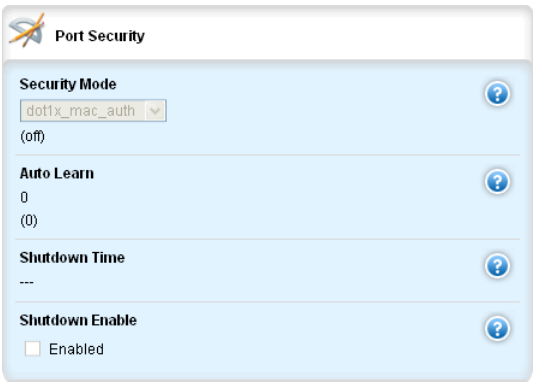

Figure 24.6. Port Security form

Security Mode

Synopsis: *string - one of the following keywords { dot1x_mac_auth, dot1x, per_macaddress, off }*

Default: *off*

Enables or disables the security feature for the port. The following port access control types are available:

- Static MAC address based. With this method, authorized MAC address(es) should be configured in the static MAC address table. If some MAC addresses are not known in advance (or which port they are going to reside behind is unknown), there is still an option to configure the switch to auto-learn a certain number of MAC addresses. Once learned, they don't age out until the unit is reset or the link goes down.
- IEEE 802.1X standard authentication.
- IEEE 802.1X with MAC Authentication, also known as MAC-Authentication Bypass. With this method, the device can authenticate clients based on the client's MAC address, if IEEE 802.1X authentication times out.

Auto Learn

Synopsis: *integer*

Default:

The maximum number of MAC addresses that can be dynamically learned on the port. If there are static addresses configured on the port, the actual number of addresses allowed to be learned is this number minus the number of the static MAC addresses.

Shutdown Time

Synopsis: *integer*

How long to shut down an interface if a security violation occurs.

Shutdown Enable

Enables/disables administative shutdown if a security violation occurs.

24.2.2. 802.1X Parameters

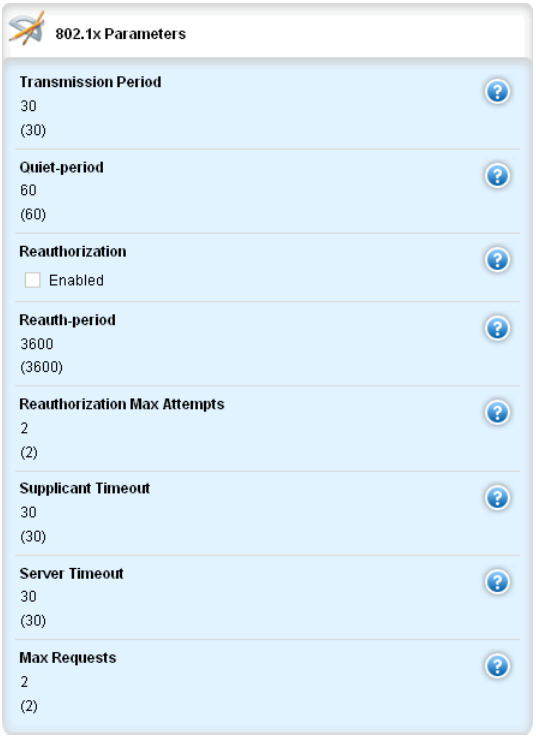

Figure 24.7. 802.1x Parameters form

Transmission Period

Synopsis: *integer*

Default: *30*

IEEE 802.1X PAE (Port Access Entity) parameters

quiet-period

Synopsis: *integer*

Default: *60*

The period of time not to attempt to acquire a supplicant after the authorization session failed.

Reauthorization

Enables or disables periodic reauthentication

reauth-period

Synopsis: *integer*

Default: *3600*

The time between successive reauthentications of the supplicant.

Reauthorization Max Attempts

Synopsis: *integer*

Default: *2*

The number of reauthentication attempts that are permitted before the port becomes unauthorized.

Supplicant Timeout

Synopsis: *integer*

Default: *30*

The time to wait for the supplicant's response to the authentication server's EAP packet.

Server Timeout

Synopsis: *integer*

Default: *30*

The time to wait for the authentication server's response to the supplicant's EAP packet.

Max Requests

Synopsis: *integer*

Default: *2*

The maximum number of times to retransmit the authentication server's EAP Request packet to the supplicant before the authentication session times out.

25. Multicast Filtering

ROX™ Multicast Filtering provides the following features:

- Support for up to 256 Multicast Groups (either static or dynamic).
- Ability to prioritize a Static Multicast Group via Class-of-Service.
- Industry standard support of IGMP (RFC 1112, RFC 2236) versions 1 and 2 in active and passive roles.
- Support of IEEE 802.1Q-2005 standard GMRP (GARP Multicast Registration protocol).
- Ability to enable or disable IGMP on a per VLAN basis.
- Multicast routers may be statically configured or dynamically recognized by IGMP.
- "Routerless" IGMP operation.

ROX™ performs Multicast Filtering using the following methods:

- Static Multicast Groups.
- Internet Group Management Protocol (IGMP) snooping.
- IEEE standard GARP Multicast Registration protocol (GMRP).

ROX™ IGMP Snooping supports multicast routers using IGMP version 2 and hosts using either IGMP version 1 or 2.

25.1. IGMP

IGMP is used by IP hosts to report their host group memberships to multicast routers. As hosts join and leave specific multicast groups, streams of traffic are directed to or withheld from that host.

The IGMP protocol operates between multicast routers and IP hosts. When an unmanaged switch is placed between multicast routers and their hosts, the multicast streams will be distributed to all ports. This may introduce significant traffic onto ports that do not require it and receive no benefit from it.

RuggedCom products with IGMP Snooping enabled will act on IGMP messages sent from the router and the host, restricting traffic streams to the appropriate LAN segments.

25.1.1. Router and Host IGMP Operation

The network shown in Figure 25.1, “IGMP Operation Example 1” provides a simple example of the use of IGMP. One “producer” IP host (P1) is generating two IP multicast streams, M1 and M2. There are four potential “consumers” of these streams, C1 through C4.

The multicast router discovers which host wishes to subscribe to which stream by sending general membership queries to each of the segments.

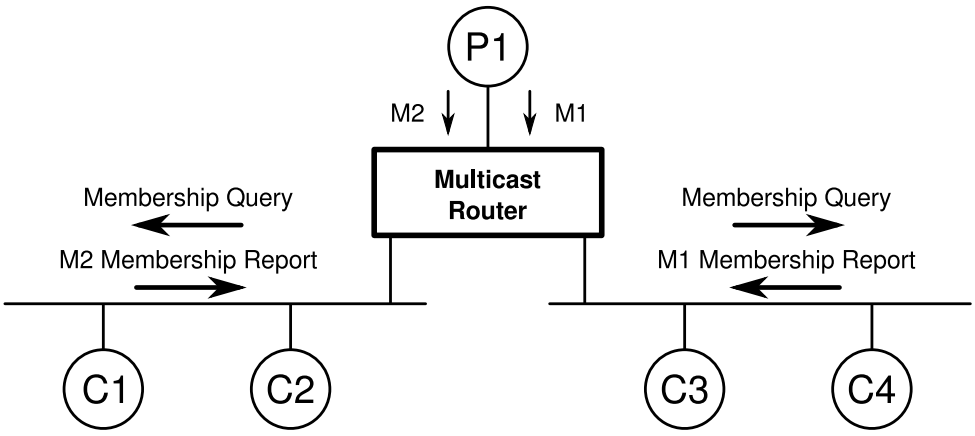

Figure 25.1. IGMP Operation Example 1

In this example, the general membership query sent to the C1-C2 segment is answered by a membership report indicating the desire to subscribe to a stream M2. The router will forward the M2 stream onto the C1-C2 segment. In a similar fashion, the router discovers that it must forward M1 onto segment C3-C4.

Membership reports are also referred to as “joins”.

A “consumer” may join any number of multicast groups, issuing a membership report for each group. When a host issues a membership report, other hosts on the same network segment that also require membership to the same group suppress their own requests, since they would be redundant. In this way, the IGMP protocol guarantees that the segment will issue only one join for each group.

The router periodically queries each of its segments in order to determine whether at least one consumer still subscribes to a given stream. If it receives no responses within a given timeout period (usually two query intervals), the router will prune the multicast stream from the given segment.

A more usual method of pruning occurs when consumers wishing to unsubscribe issue an IGMP “leave group” message. The router will immediately issue a group-specific membership query to determine whether there are any remaining subscribers of that group on the segment. After the last consumer of a group has un-subscribed, the router will prune the multicast stream from the given segment.

25.1.2. Switch IGMP Operation

The IGMP Snooping feature provides a means for switches to snoop (i.e. watch) the operation of routers, respond with joins/leaves on the behalf of consumer ports and to prune multicast streams accordingly.

There are two modes of IGMP that the switch can be configured to assume - active and passive.

Active Mode

ROX™ IGMP supports a “routerless” mode of operation.

When such a switch is used without a multicast router, it is able to function as if it is a multicast router sending IGMP general queries.

Passive Mode

When such a switch is used in a network with a multicast router, it can be configured to run Passive IGMP. This mode prevents the switch from sending the queries that can confuse the router causing it to stop issuing IGMP queries.

A switch running in passive mode requires the presence of a multicast router or it will not be able to forward multicast streams at all

If no multicast routers are present, at least one IGMP Snooping switch must be configured for Active IGMP mode to make IGMP functional.

IGMP Snooping Rules

- When a multicast source starts multicasting, the traffic stream will be immediately blocked on segments from which joins have not been received.
- The switch will always forward all multicast traffic to the ports where multicast routers are attached unless configured otherwise.
- Packets with a destination IP multicast address in the 224.0.0.X range which are not IGMP are always forwarded to all ports. This behavior is based on the fact that many systems do not send joins for IP multicast addresses in this range while still listening to such packets.
- The switch implements “proxy-reporting”, i.e. membership reports received from downstream are summarized and used by the switch to issue its own reports.
- The switch will only send IGMP membership reports out of those ports where multicast routers are attached because sending membership reports to hosts could result in unintentionally preventing a host from joining a specific group.
- Multicast routers use IGMP to elect a master router known as the querier – the one with the lowest IP address is elected to be the querier, all other routers become non-queriers, participating only forward multicast traffic. Switches running in Active IGMP mode participate in the querier election like multicast routers.
- When the querier election process is complete, the switch simply relays IGMP queries received from the querier.
- When sending IGMP packets, the switch uses its own IP address, if it has one, for the VLAN on which packets are sent, or an address of 0.0.0.0, if it doesn’t have an assigned IP address.

IGMP Snooping switches perform multicast pruning using a multicast frames’ destination MAC multicast address which depends on the group IP multicast address. IP address W.X.Y.Z corresponds to MAC address 01-00-5E-XX-YY-ZZ where XX is the lower 7 bits of X and YY and ZZ are simply Y and Z coded in hexadecimal.

One can note that IP multicast addresses such as 224.1.1.1 and 225.1.1.1 will both map onto the same MAC address 01-00-5E-01-01-01. This is indeed a problem for which the IETF Network Working Group currently has offered no solution. Users are advised to be aware of and avoid this problem.

IGMP and RSTP

An RSTP change of topology can render the routes selected to carry multicast traffic as incorrect. This results in lost multicast traffic.

If RSTP detects change in the network topology, IGMP will take some actions to avoid loss of multicast connectivity and reduce network convergence time:

- The switch will immediately issue IGMP queries (if in IGMP Active mode) to obtain potential new group membership information.
- The switch can be configured to flood multicast streams temporarily out of all ports that are not configured as RSTP Edge Ports.

25.1.3. Combined Router and Switch IGMP Operation

This section describes the additional challenges of multiple routers, VLAN support and switching.

Producer P1 resides on VLAN 2 while P2 resides on VLAN 3. Consumer C1 resides on both VLANs whereas C2 and C3 reside on VLANs 3 and 2, respectively. Router 2 resides on VLAN 2, presumably to forward multicast traffic to a remote network or act as a source of multicast traffic itself.

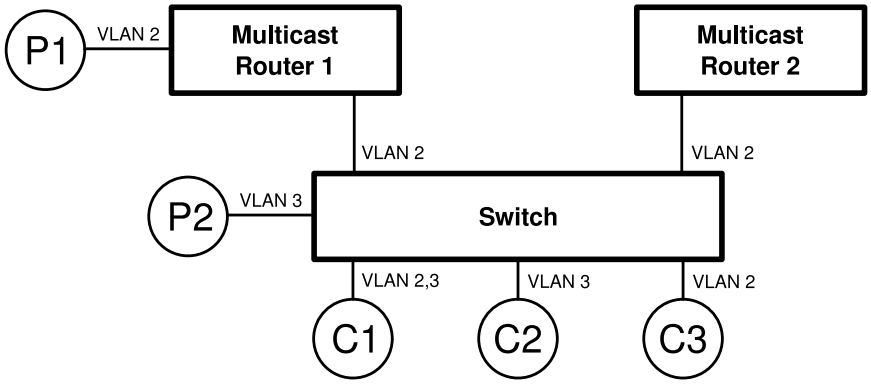

Figure 25.2. IGMP Operation Example 2

In this example, we will assume that all the devices agree that router 1 is the querier for VLAN 2 and router 2 is simply a non-querier. In this case, the switch will periodically receive queries from router 1 and, thus, maintain the information concerning which of its ports links to the multicast router. However, the switch port that links to router 2 must be manually configured as a “router port”. Otherwise, the switch will send neither multicast streams nor joins/leaves to router 2.

Note that VLAN 3 does not have an external multicast router. The switch should be configured to operate in its “routerless” mode and issue general membership queries as if it is the router.

Processing Joins

If host C1 desires to subscribe to the multicast streams for both P1 and P2, it will generate two joins. The join from C1 on VLAN 2 will cause the switch to immediately initiate its own join to multicast router 1 (and to issue its own join as a response to queries).

The join from C1 for VLAN 3 will cause the switch to immediately begin forwarding multicast traffic from P2 to C2.

Processing Leaves

When host C1 decides to leave a multicast group, it will issue a leave request to the switch. The switch will poll the port to determine if C1 is the last member of the group on that port. If C1 is the last (or only) member, the group will immediately be pruned from the port.

Should host C1 leave the multicast group without issuing a leave group message and then fail to respond to a general membership query, the switch will stop forwarding traffic after two queries.

When the last port in a multicast group leaves the group (or is aged-out), the switch will issue an IGMP leave report to the router.

25.2. GMRP (GARP Multicast Registration Protocol)

The GARP Multicast Registration Protocol (GMRP) is an application of the Generic Attribute Registration Protocol (GARP) that provides a mechanism at Layer 2 for managing multicast group membership in a bridged Layer 2 network. It allows Ethernet switches and end stations to register and unregister

membership in multicast groups with other switches on a LAN, and for that information to be disseminated to all switches in the LAN that support Extended Filtering Services.

GMRP is an industry-standard protocol first defined in IEEE 802.1D-1998 and extended in IEEE 802.1Q-2005. GARP was defined in IEEE 802.1D-1998 and updated in 802.1D-2004. Note that GMRP provides similar functionality at Layer 2 to that which IGMP, described in the preceding sections, provides at Layer 3.

Joining a Multicast Group)

In order to join a multicast group, an end station transmits a GMRP “join” message. The switch that receives the “join” message adds the port through which the message was received to the multicast group specified in the message. It then propagates the “join” message to all other hosts in the VLAN, one of which is expected to be the multicast source.

When a switch transmits GMRP updates (from GMRP-enabled ports), all of the multicast groups known to the switch, whether configured manually or learned dynamically through GMRP, are advertised to the rest of network.

As long as one host on the Layer 2 network has registered for a given multicast group, traffic from the corresponding multicast source will be carried on the network. Traffic multicast by the source is only forwarded by each switch in the network to those ports from which it has received join messages for the multicast group.

Leaving a Multicast Group

Periodically, the switch sends GMRP queries in the form of a “leave all” message. If a host (either a switch or an end station) wishes to remain in a multicast group, it reasserts its group membership by responding with an appropriate “join” request. Otherwise, it can either respond with a “leave” message or simply not respond at all. If the switch receives a “leave” message or receives no response from the host for a timeout period, the switch removes the host from the multicast group.

GMRP Protocol Notes

Since GMRP is an application of GARP, transactions take place using the GARP protocol. GMRP defines the following two Attribute Types:

- The *Group Attribute Type*, used to identify the values of group MAC addresses
- The *Service Requirement Attribute Type*, used to identify service requirements for the group

Service Requirement Attributes are used to change the receiving port’s multicast filtering behavior to one of the following:

- Forward All Multicast group traffic in the VLAN, or
- Forward All Unknown Traffic (Multicast Groups) for which there are no members registered in the device in a VLAN

If GMRP is disabled on the RuggedBackbone™, then GMRP PDUs received by the RX1500 will be forwarded like any other traffic; but if GMRP is enabled on at least one of the ports, then GMRP packets will be processed by the RX1500, and not forwarded.

25.2.1. GMRP Example

In the example depicted in Figure 25.3, “Example using GMRP”, there are two multicast sources, S1 and S2, multicasting to Multicast Groups 1 and 2, respectively. A network of five switches, including one core Switch, B, connects the sources to two hosts, H1 and H2, which receive the multicast streams from S1 and S2, respectively.

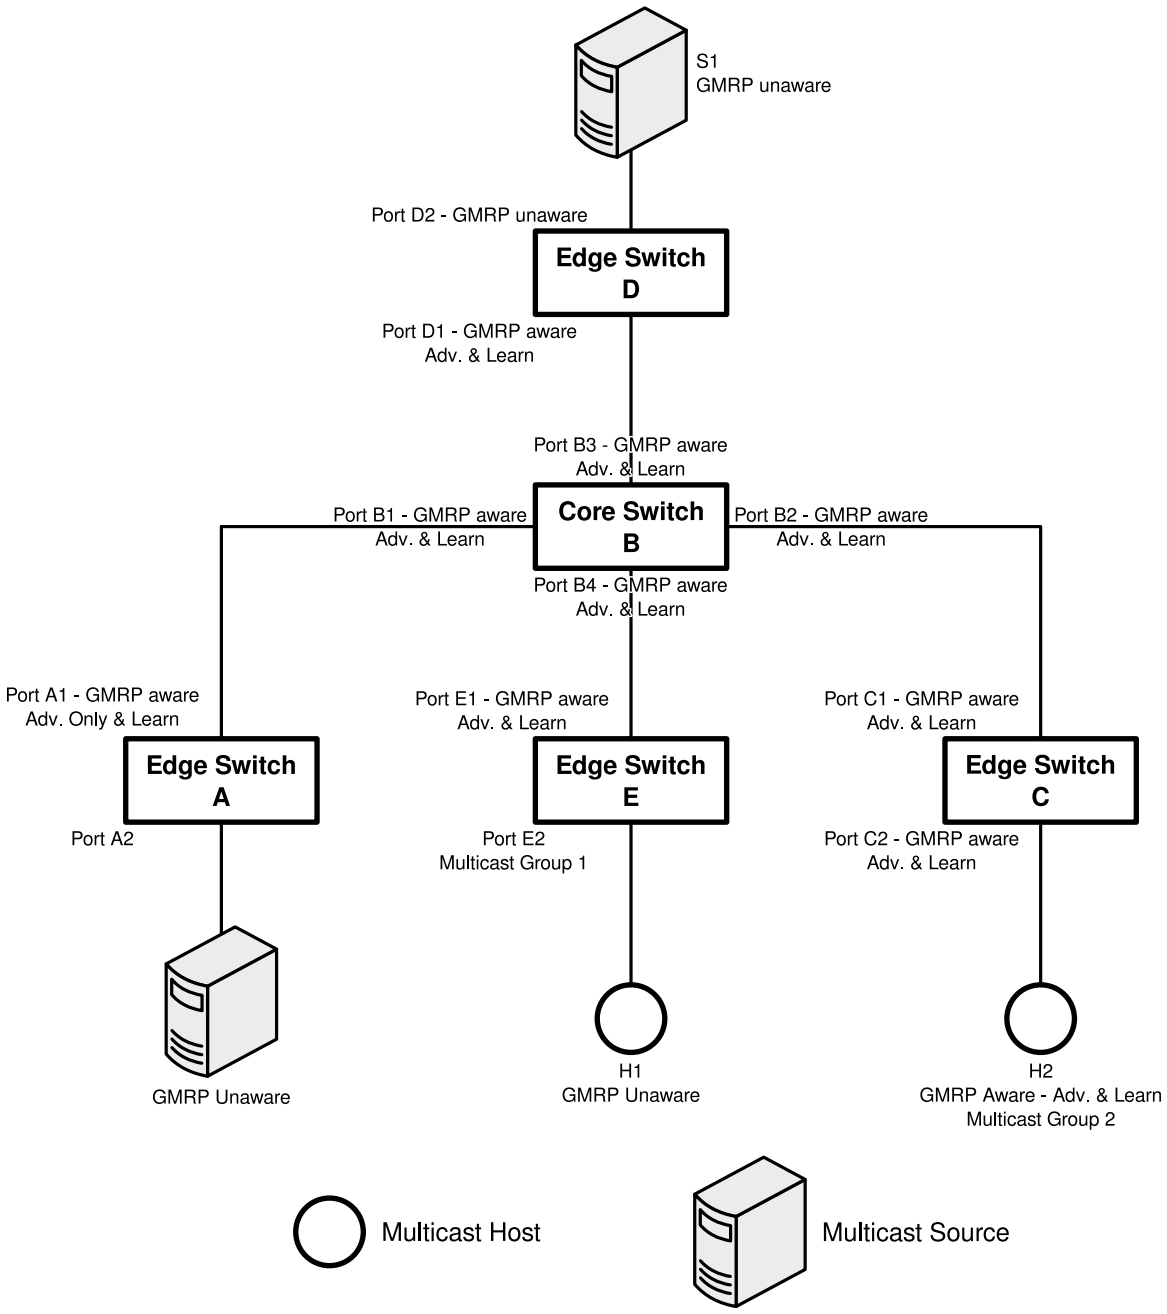

Figure 25.3. Example using GMRP

Joining the Multicast Groups:

The sequence of events surrounding the establishment of membership for the two Multicast Groups on the example network is as follows:

- Host H1 is GMRP unaware but needs to see traffic for Multicast Group 1. Port E2 on Switch E, therefore, is statically configured to forward traffic for Multicast Group 1.
- Switch E advertises membership in Multicast Group 1 to the network through Port E1, making Port B4 on Switch B a member of Multicast Group 1.

- Switch B propagates the “join” message, causing Port D1 on Switch D to become a member of Multicast Group 1. Note that ports A1 and C1 also become members.
- Host H2 is GMRP-aware and sends a “join” request for Multicast Group 2 to Port C2, which thereby becomes a member of Group 2.
- Switch C propagates the “join” message, causing Port B2 on Switch B and Port A1 on Switch A to become members of Multicast Group 2. Note that ports D1 and E1 also become members.

Multicast Traffic on the Network

Once GMRP-based registration has propagated through the network as described above, multicasts from S1 and S2 can reach their destinations, as described in the following:

- Source S1 transmits multicast traffic to Port D2 which is forwarded via Port D1, which has previously become a member of Multicast Group 1.
- Switch B forwards the Group 1 multicast via Port B4 towards Switch E.
- Switch E forwards the Group 1 multicast via Port E2, which has been statically configured for membership in Multicast Group 1.
- Host H1, connected to Port E2, thus receives the Group 1 multicast.
- Source S2 transmits multicast traffic to Port A2, which is then forwarded via port A1, which has previously become a member of Multicast Group 2.
- Switch B forwards the Group 2 multicast via Port B2 towards Switch C.
- Switch C forwards the Group 2 multicast via Port C2, which has previously become a member of Group 2.
- Ultimately, Host H2, connected to Port C2, receives the Group 2 multicast.

25.3. Multicast Filtering Configuration and Status

The Multicast Filtering menu is accessible from the main menu under *switch*. The path to this menu is *switch/mcast-filtering*.

Figure 25.4. Multicast Filtering menu

25.3.1. Configuring IGMP Parameters

Note that the activation of IGMP on a per-VLAN basis is configured using Static VLANs.

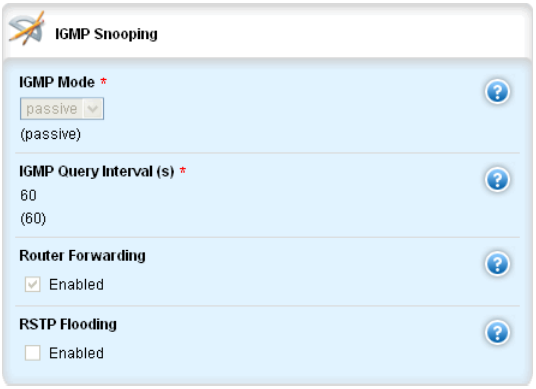

Figure 25.5. IGMP Snooping Parameters form

The path to the IGMP Snooping forms and the Router Ports table is *switch/mcast-filtering/igmp-snooping*.

IGMP Mode

Synopsis: *string - one of the following keywords { passive, active }*

Default: *passive*

Specifies the IGMP mode:

- PASSIVE : the switch passively snoops IGMP traffic and never sends IGMP queries.
- ACTIVE : the switch generates IGMP queries, if no queries from a better candidate for being the querier are detected for a while.

IGMP Query Interval (s)

Synopsis: *integer*

Default: *60*

The time interval between IGMP queries generated by the switch. NOTE: This parameter also affects the Group Membership Interval (i.e. the group subscriber aging time), therefore, it takes effect even in PASSIVE mode.

Router Forwarding

Synopsis: *boolean*

Default: *true*

Whether or not multicast streams will always be forwarded to multicast routers.

RSTP Flooding

Whether or not multicast streams will be flooded out of all RSTP non-edge ports upon detection of a topology change. Such flooding is desirable, if multicast stream delivery must be guaranteed without interruption.

Figure 25.6. Router Ports table

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

25.3.2. Configuring Static Multicast Groups

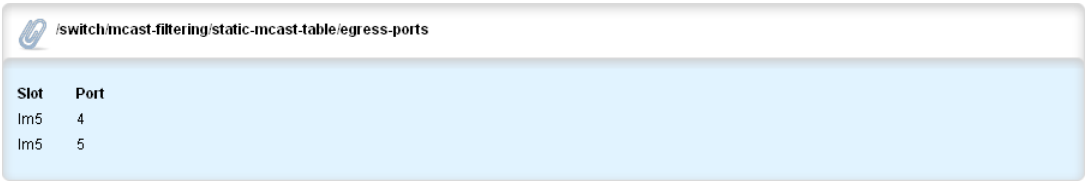

Figure 25.7. Egress Ports table

If data is configured, display the Egress Ports table by navigating to *switch/mcast-filtering/static-mcast-table* and then clicking on one of the linked submenus.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

Figure 25.8. Static Multicast Summary table

If data is configured, the path to the Static Multicast Summary table will be *switch/mcast-filtering/static-mcast-table*.

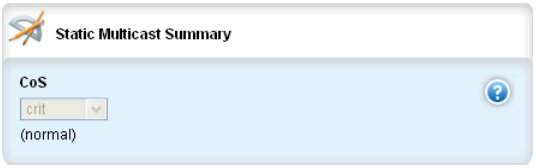

Figure 25.9. Static Multicast Summary form

If data is configured, the path to the Static Multicast Summary form will be *switch/mcast-filtering/static-mcast-table* and then clicking on one of the linked submenus that follow.

VLAN ID

Synopsis: *integer*

The VLAN Identifier of the VLAN upon which the multicast group operates.

MAC Address

Synopsis: *Multicast Ethernet MAC address in colon-separated hexadecimal notation*

The Multicast group MAC address in the form 01:xx:xx:xx:xx:xx.

CoS

Synopsis: *string - one of the following keywords { crit, high, medium, normal }*

Default: *normal*

The Class Of Service that is assigned to the multicast group frames.

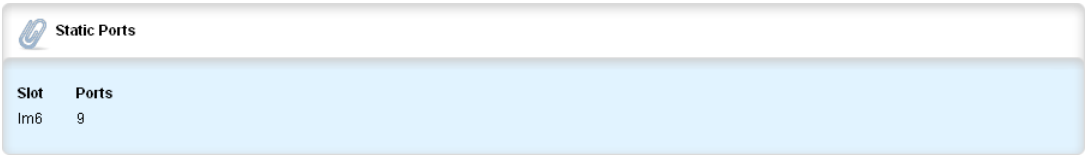

Figure 25.10. Static Ports table

If data is configured, the path to this menu will be *switch/mcast-filtering/mcast-group-summary*, then clicking on one of the linked submenus and then clicking on *static-ports*.

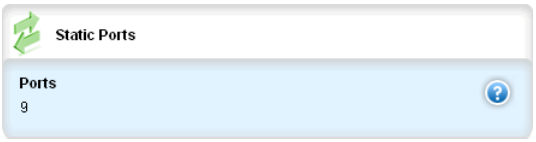

Figure 25.11. Static Ports form

If data is configured, the path to this menu will be *switch/mcast-filtering/mcast-group-summary* then clicking on one of the linked submenus, then clicking on *static-ports* and then on a linked submenu.

Static-ports are egress ports that have been assigned to a particular multicast MAC address in the static-mcast-table menu.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Ports

Synopsis: *A string*

The selected ports on the module installed in the indicated slot.

25.3.2.1. Multicast Group Summary

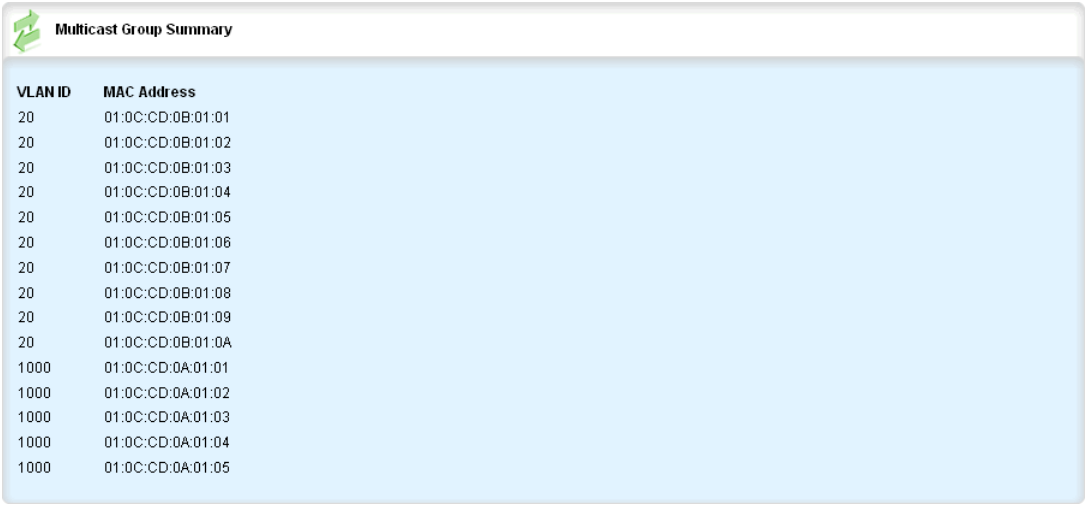

Figure 25.12. Multicast Group Summary table

The path to this table is *switch/mcast-filtering/mcast-group-summary*.

VLAN ID

Synopsis: *integer*

The VLAN Identifier of the VLAN upon which the multicast group operates.

MAC Address

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The multicast group MAC address.

25.3.2.2. Viewing IP Multicast Groups

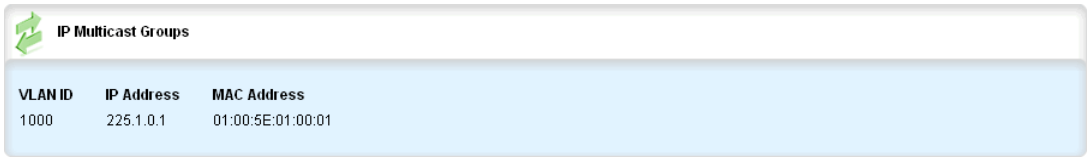

Figure 25.13. IP Multicast Groups table

The IP Multicast Groups table allows you to view IP Multicast Groups.

The path to this table is *switch/mcast-filtering/ip-mcast-groups*.

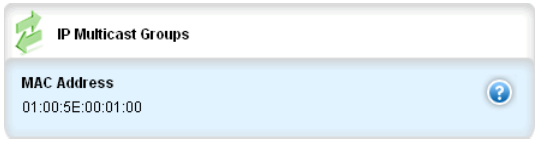

Figure 25.14. IP Multicast Groups form

The path to this form is *switch/mcast-filtering/ip-mcast-groups* and then clicking on one of the linked submenus that follow.

VLAN ID

Synopsis: *integer*

The VLAN Identifier of the VLAN upon which the multicast group operates.

IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

The multicast group IP address

MAC Address

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The multicast MAC address corresponding to the group multicast IP address.

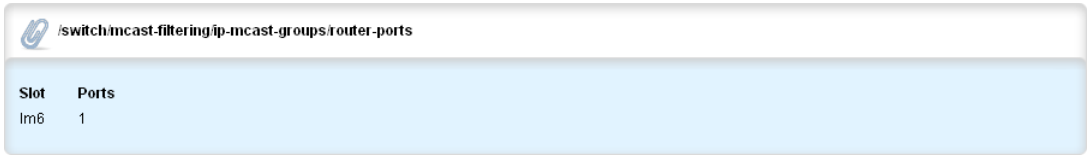

Figure 25.15. Router-Ports table

The path to this table is *switch/mcast-filtering/ip-mcast-groups*, then clicking on one of the linked submenus that follow, and then clicking on *router-ports*.

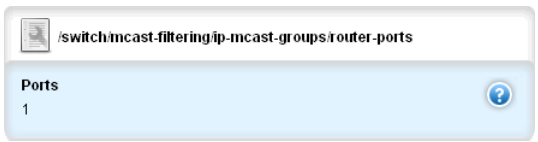

Figure 25.16. Router-Ports form

The path to this form is *switch/mcast-filtering/ip-mcast-groups*, then clicking on one of the linked submenus that follow, then on *router-ports* and then on a linked submenu.

All ports that have been manually configured or dynamically discovered (by observing router specific traffic) as ports that link to multicast routers.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Ports

Synopsis: *A string*

The selected ports on the module installed in the indicated slot.

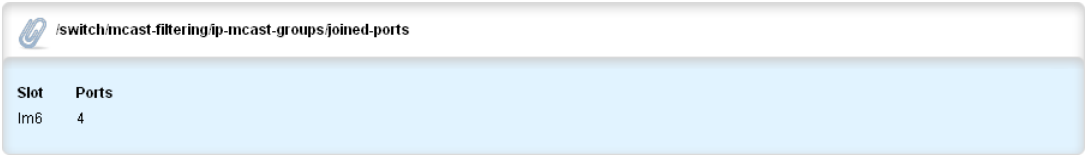

Figure 25.17. Joined-Ports table

The path to this table is *switch/mcast-filtering/ip-mcast-groups*, then clicking on one of the linked submenus that follow, and then clicking on *joined-ports*.

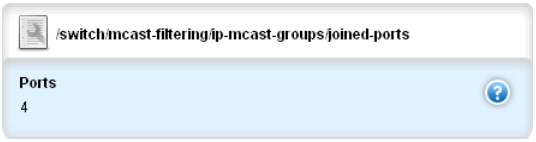

Figure 25.18. Joined-Ports form

The path to this form is *switch/mcast-filtering/ip-mcast-groups*, then clicking on one of the linked submenus that follow, then on *joined-ports* and then on a linked submenu.

All ports that subscribed to the multicast group traffic.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Ports

Synopsis: *A string*

The selected ports on the module installed in the indicated slot.

25.3.3. Configuring GMRP

Figure 25.19. GMRP form

The GMRP Form appears on the same screen as the *Multicast Filtering* menu.

Enabled

Synopsis: *boolean*

Default: *false*

GMRP Enable

RSTP Flooding

Whether or not multicast streams will be flooded out of all RSTP non-edge ports upon detection of a topology change. Such flooding is desirable, if multicast stream delivery must be guaranteed without interruption.

Leave Timer (ms)

Synopsis: *integer*

Default: *4000*

The time in milliseconds to wait after issuing Leave or LeaveAll before removing registered multicast groups. If Join messages for specific addresses are received before this timer expires, the addresses will be kept registered.

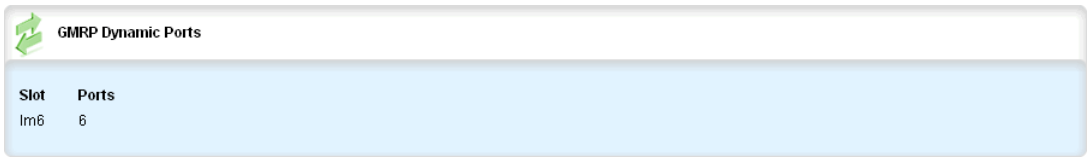

Figure 25.20. GMRP Dynamic Ports table

The path to this menu is *switch/mcast-filtering/mcast-group-summary*, then clicking on one of the linked submenus and then clicking on *gmrp-dynamic-ports*.

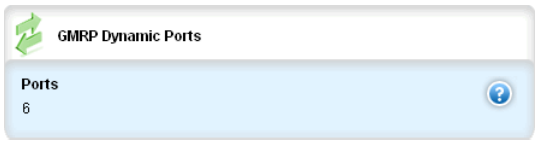

Figure 25.21. GMRP Dynamic Ports form

The path to this menu is *switch/mcast-filtering/mcast-group-summary*, then clicking on one of the linked submenus, then on *gmrp-dynamic-ports* and finally, clicking on a linked submenu that follows.

GMRP Dynamic Ports are ports that joined this group dynamically through a GMRP Application and to which the multicast group traffic is forwarded.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Ports

Synopsis: *A string*

The selected ports on the module installed in the indicated slot.

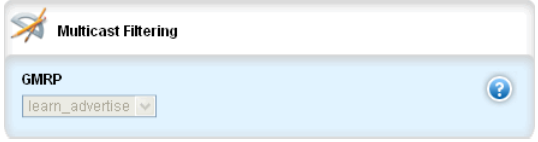

Figure 25.22. Multicast Filtering form

The Multicast Filtering form can be accessed in two locations: *interface/switch* and then clicking on a submenu (for example, lm1/1) or *interface/trunks* and then clicking on a submenu (for example, 1).

GMRP

Synopsis: *string - one of the following keywords { learn_advertise, advertise_only }*

GMRP (GARP Multicast Registration Protocol) operation on the port. There are several GMRP operation modes:

- DISABLED : the port is not capable of any GMRP processing.
- ADVERTISE ONLY : the port will declare all MCAST addresses existing in the switch (configured or learned) but will not learn any MCAST addresses.
- ADVERTISE and LEARN : the port will declare all MCAST Addresses existing in the switch (configured or learned) and can dynamically learn MCAST addresses.

25.4. Troubleshooting

Problem One

When I start a multicast traffic feed, it is always distributed to all members of the VLAN.

Is IGMP enabled for the VLAN? Multicasts will be distributed to all members of the VLAN unless IGMP is enabled.

Problem Two

Computers on my switch receive the multicast traffic just fine, but I can't get the stream through a connected router.

Is the port used to connect the router included in the Router Ports list?

To determine whether the multicast stream is being delivered to the router, run the Ethernet Statistics menu *View Ethernet Statistics* command. Verify that the traffic count transmitted to the router is the same as the traffic count received from the multicasting source.

Problem Three

The video stream at one of my end stations is of pretty poor quality.

Video serving is a resource-intensive application. Because it uses isochronous workload, data must be fed at a prescribed rate or end users will see glitches in the video. Networks that carry data from the server to the client must be engineered to handle this heavy, isochronous workload.

Video streams can consume large amounts of bandwidth. Features and capacity of both server and network (including routers, bridges, switches, and interfaces) impact the streams.

You should not exceed 60% of the maximum interface bandwidth. For example, if using a 10 Mbps Ethernet, you should run a single multicasting source at no more than 6 Mbps, or two sources at 3 Mbps.

Router ports will carry the traffic of all multicast groups, so it is especially important to consider these ports in your design.

Note that multicasting will definitely introduce latency in all traffic on the network. Plan your network carefully in order to account for capacity and latency concerns.

Problem Four

Multicast streams of some groups are not forwarded properly. Some segments without subscribers receive the traffic while some segments with subscribers don't.

Ensure that you do not have a situation where different multicast groups have multicast IP addresses that map to the same multicast MAC address. The switch forwarding operation is MAC address-based and will not work properly for several groups mapping to the same MAC address.

Problem Five

Computers on my switch issue join requests but don't receive multicast streams from a router.

Is your multicast router running IGMP version 2? It must run IGMP version 2 in order for IGMP Snooping to operate properly.

Problem Six

I connect or disconnect some switch ports and multicast goes everywhere. Is IGMP broken?

No, it may be a proper switch behavior. When the switch detects a change in the network topology through RSTP, it acts to avoid loss of multicast traffic – if configured to do so, it starts forwarding all multicast traffic to all ports that are not RSTP Edge ports (because they may potentially link to routers). This may result in some undesired flooding of multicast traffic (which will stop after a few minutes), however, it guarantees that all devices interested in the traffic will keep receiving it without a break. Note that the same behavior will be observed when the switch resets or when IGMP Snooping is being disabled for the VLAN.

26. Classes Of Service

ROX™ CoS provides the following features:

- Support for 4 Classes of Service
- Ability to prioritize traffic by ingress port.
- Ability to prioritize traffic by the priority field in 802.1Q tags.
- Ability to prioritize traffic based on its source or destination MAC address.
- Ability to prioritize traffic by the TOS field in the IP header.

26.1. CoS Operation

CoS provides the ability to expedite the transmission of certain frames and port traffic over others.

The CoS of a frame can take on one of four values: Normal, Medium, High or Critical.

The default policies of the switch enforce a Normal CoS for all traffic.

Use the highest supported CoS with caution, as it is always used by the switch for handling network management traffic such as RSTP BPDUs.

If this CoS is used for regular network traffic, upon traffic bursts, it may result in loss of some network management frames which in its turn may result in loss of connectivity over the network.

The CoS feature has two main phases - inspection and forwarding:

26.1.1. Inspection Phase

In the inspection phase, the CoS priority of a received frame is determined from:

- A specific CoS based upon the source and destination MAC address (as set in the Static MAC Address Table).
- The priority field in 802.1Q tags.
- The Differentiated Services Code Point (DSCP) component of the Type Of Service (TOS) field, if the frame is IP.
- The default CoS for the port.

Note that a frame's CoS will be determined once the first examined parameter is found in the frame.

Received frames are first examined to determine if their destination or source MAC address is found in the Static MAC Address Table. If yes, the CoS configured for the static MAC address is used. If neither destination or source MAC address is in the Static MAC Address Table, the frame is then examined for 802.1Q tags and the priority field is mapped to a CoS. If a tag is not present, the frame is examined to determine if it is an IP frame. If the frame is IP and inspecting TOS is enabled, the CoS is determined from the DSCP field. If the frame is not IP or inspecting TOS is disabled, the default CoS for the port is used.

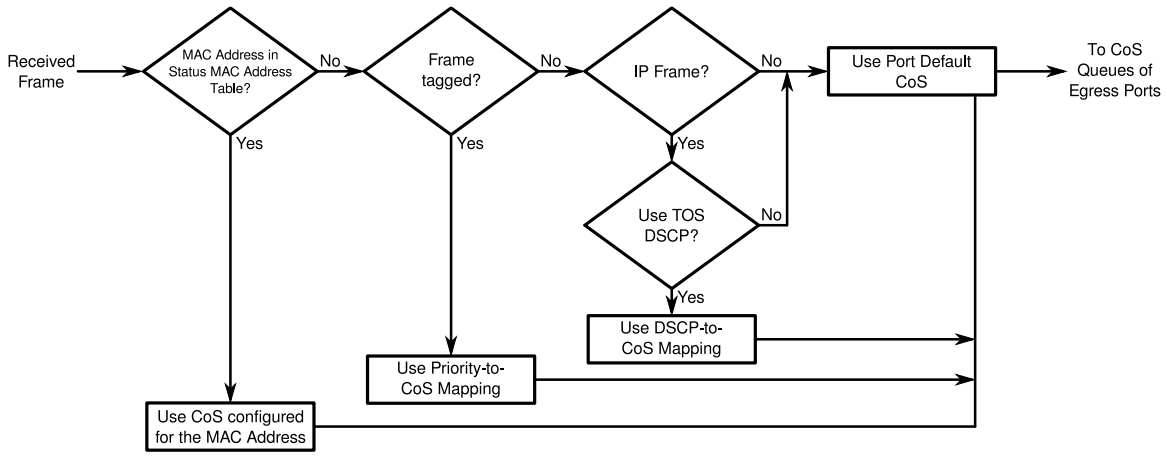

Figure 26.1. Determining The CoS Of A Received Frame

After inspection, the frame is forwarded to the egress port for transmission.

26.1.2. Forwarding Phase

The inspection phase results in the CoS of individual frames being determined. When these frames are forwarded to the egress port, they are collected into one of the priority queues according to the CoS assigned to each frame.

CoS weighting selects the degree of preferential treatment that is attached to different priority queues. The ratio of the number of higher CoS to lower CoS frames transmitted can be programmed. If desired, the user can program lower CoS frames are to be transmitted only after all higher CoS frames have been serviced.

26.2. CoS Configuration

The class-of-service menu is accessible from the main menu under *switch*. The path to this menu is *switch/class-of-service*.

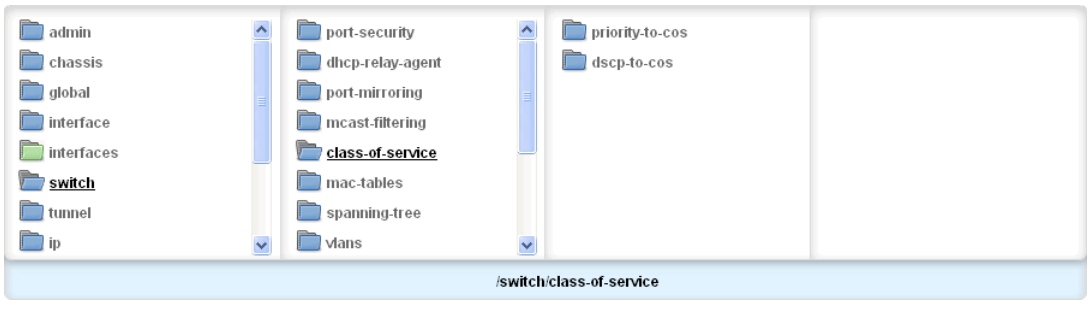

Figure 26.2. Class-of-service menu

26.2.1. Global CoS Parameters

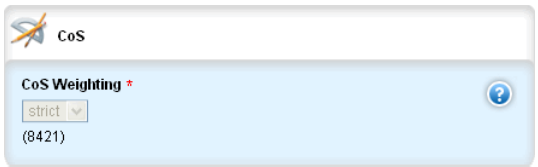

Figure 26.3. CoS form

The CoS form appears on the same screen as the *Class-of-service* menu.

CoS Weighting

Synopsis: *string - one of the following keywords { strict, 8421 }*

Default: *8421*

During traffic bursts, frames queued in the switch pending transmission on a port may have different CoS priorities. This parameter specifies the weighting algorithm for transmitting different priority CoS frames.

26.2.2. Priority to CoS Mapping

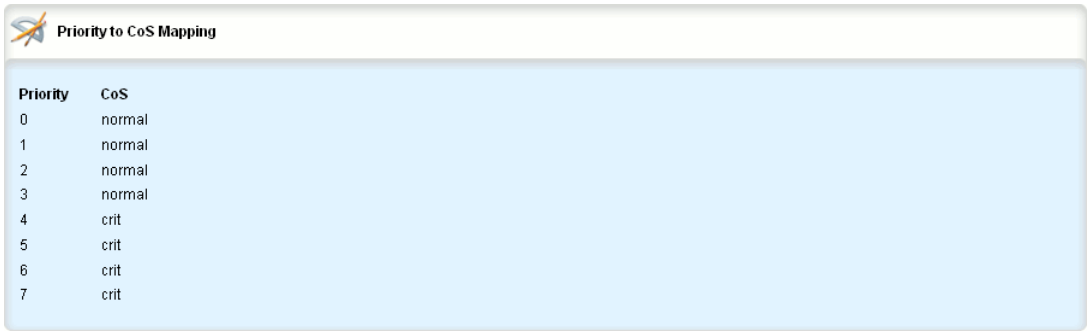

Figure 26.4. Priority to CoS Mapping table

The path to the Priority to CoS Mapping table is *switch/class-of-service/priority-to-cos*. This table shows the mapping of each IEEE 802.1p priority value to the Class of Service switch.

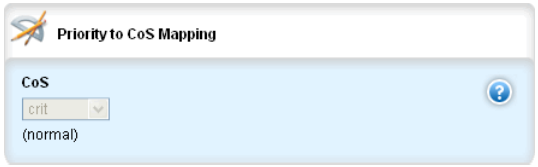

Figure 26.5. Priority to CoS Mapping form

The path to the Priority to CoS Mapping forms is *switch/class-of-service/priority-to-cos/1*. Note that any of the linked submenus from 0 to 7 can be clicked to get to the forms.

Priority

Synopsis: *integer*

The value of the IEEE 802.1p priority.

CoS

Synopsis: *string - one of the following keywords { crit, high, medium, normal }*

Default: *normal*

The CoS assigned to received tagged frames with the specified IEEE 802.1p priority value.

26.2.3. DSCP to CoS Mapping

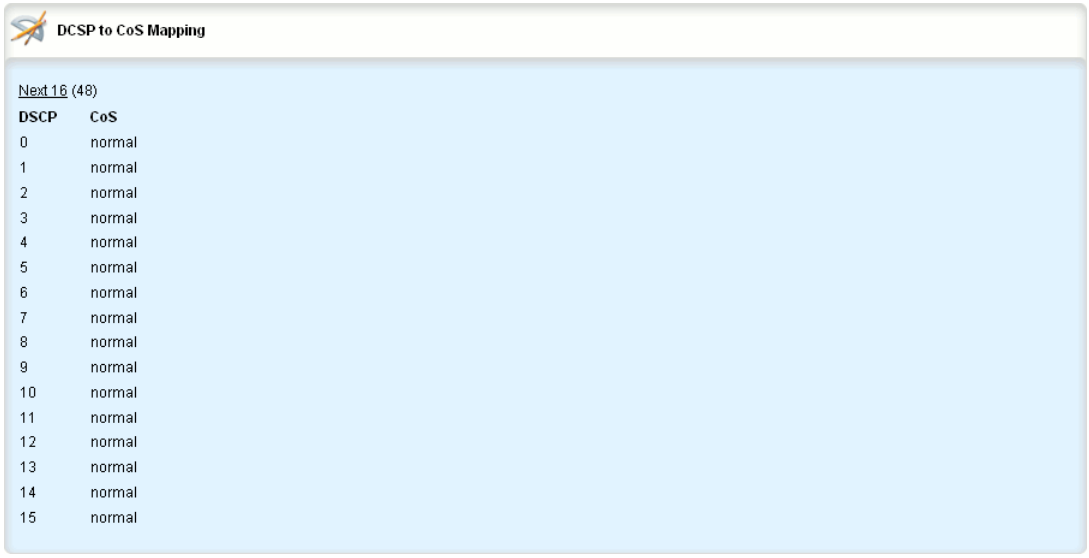

Figure 26.6. TOS DSCP to CoS Mapping table

The path to the TOS DSCP table is *switch/class-of-service/dcsp-to-cos-mapping*.

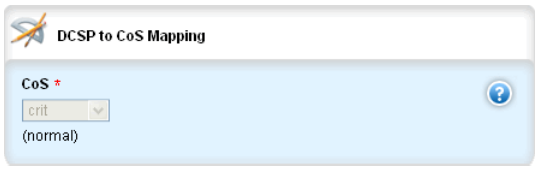

Figure 26.7. TOS DSCP to CoS Mapping form

The path to the TOS DSCP to CoS Mapping forms is *switch/class-of-service/dscp-to-cos/{number}*. TOS DSCP to CoS Mapping maps each Differentiated Services Code Point (DSCP) in the Type-Of-Service (TOS) field in the headers of the received IP packets to the Class of Service switch.

DSCP

Synopsis: *unsigned byte integer*

Differentiated Services Code Point (DSCP): a value of the 6 bit DiffServ field in the Type-Of-Service (TOS) field of the IP header.

CoS

Synopsis: *string - one of the following keywords { crit, high, medium, normal }*

Default: *normal*

The Class of Service assigned to received frames with the specified DSCP.

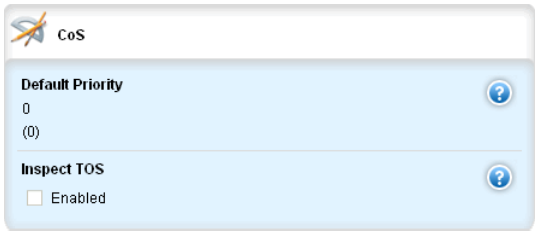

Figure 26.8. CoS form

The CoS form can be accessed in two locations: *interface/switch/{line module}/* or *interface/trunks/{number}*.

Default Priority

Synopsis: *integer*

Default:

The priority of frames received on this port that are not prioritized based on the frame's contents (e.g. the priority field in the VLAN tag, DiffServ field in the IP header, prioritized MAC address).

Inspect TOS

Enables or disables parsing of the Type-of-Service (ToS) field in the IP header of the received frames to determine what Class of Service (CoS) they should be assigned. When ToS parsing is enabled the switch will use the differentiated services bits in the TOS field.

27. MAC Address Tables

ROX™ MAC address table management provides the following features:

- Viewing learned MAC addresses.
- Configuring the switch's MAC Address Aging Time.
- Configuring static MAC addresses.
- Purging MAC Address entries.

The *MAC Address Tables (mac-tables)* menu is is accessible from the main menu under *switch/mac-tables*.

Figure 27.1. MAC Tables menu

1. Viewing MAC Addresses

To display the MAC Address table, navigate to *switch/mac-tables/mac-table*.

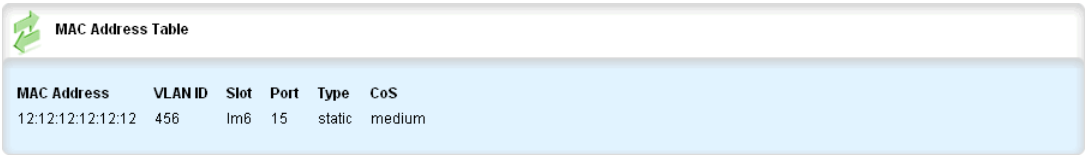

Figure 27.2. MAC Address table

To display the MAC address form, navigate to *switch/mac-tables/mac-table/mac-tables/{mac address}*.

Figure 27.3. Mac Address form

MAC Address

Synopsis: *Unicast Ethernet MAC address in colon-separated hexadecimal notation*

The MAC address learned by the switch.

VLAN ID

Synopsis: *integer*

The VLAN Identifier of the VLAN upon which the MAC address operates.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The slot containing the module including the port.

Port

Synopsis: *integer*

The port on which the MAC address has been learned.

Type

Synopsis: *string - one of the following keywords { dynamic, static }*

How the MAC address has been learned by the switch:

- STATIC : the address has been learned as a result of Static MAC Address Table configuration or some other management activity and cannot be automatically unlearned or relearned by the switch.
- DYNAMIC : the address has been automatically learned by the switch and can be automatically unlearned.

CoS

Synopsis: *string - one of the following keywords { crit, high, medium, normal }*

The Class Of Service that is assigned to frames carrying this address as source or destination address

2. Configuring MAC Address Learning Options

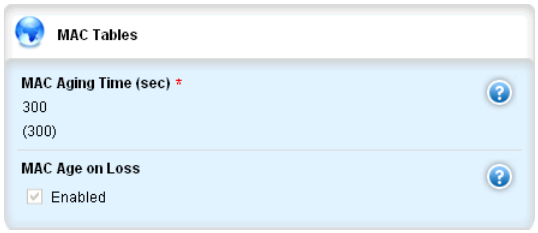

Figure 27.4. MAC Tables form

The MAC Tables form is accessible under *switch/mac-tables*.

MAC Aging Time (sec)

Synopsis: *integer*

Default: *300*

The time a learned MAC address is held before being aged out.

MAC Age on Loss

Synopsis: *boolean*

Default: *true*

When link failure (and potentially a topology change) occurs, the switch may have some MAC addresses previously learned on the failed port. As long as those addresses are not aged-out, the switch will still be forwarding traffic to that port, thus preventing that traffic from reaching its destination via the new network topology. This parameter allows the aging-out of all MAC addresses learned on a failed port immediately upon link failure detection.

3. Configuring The Static MAC Address Table

Static MAC addresses are usually configured when the user wishes to enforce port security (if supported).

Static MAC addresses must also be configured for devices that are able to receive but not able to transmit frames.

Prioritized MAC addresses are configured when traffic to or from a specific device on a LAN segment is to be assigned a higher CoS priority than other devices on that LAN segment.

To add a MAC address, go to the static-mac-table menu and enter the Edit Private mode. Click on Add Static-Mac and then use the Key Settings form to add a new MAC address. MAC Address and VLAN ID are the keys. Enter other relevant parameters in the Static MAC Address Parameters form.

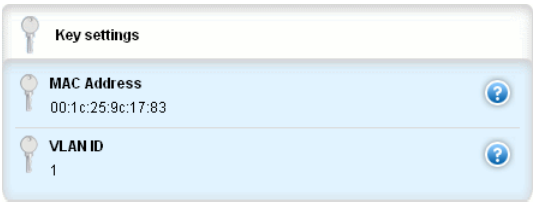

Figure 27.5. Key Settings

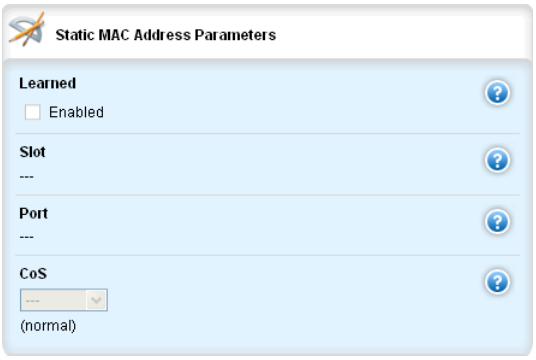

Figure 27.6. Static MAC Address Parameters form

If MAC addresses have been configured, Static MAC Address Parameters forms will appear under the static-mac-table menu. To access the forms, navigate to *switch/mac-tables/static-mac-table/{mac address}*.

Once a statically entered MAC address is discovered on the network, additional information about it will be visible via the Static MAC Address Table (*static-mac-table*):

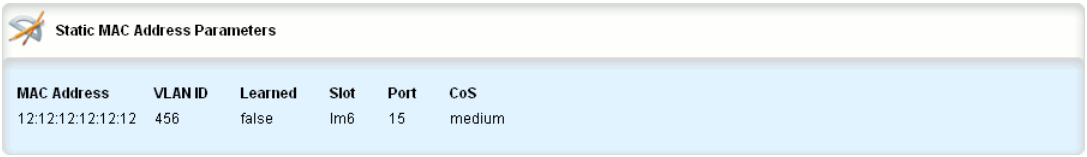

Figure 27.7. Static MAC Address Parameters table

To display the Static MAC Address Parameters table, navigate to *switch/mac-tables-static-mac-table*.

MAC Address

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

A MAC address that is to be statically configured.

VLAN ID

Synopsis: *integer*

The VLAN Identifier of the VLAN upon which the MAC address operates.

learned

If set, the system will auto-learn the port upon which the device with this address is located.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

CoS

Synopsis: *string - one of the following keywords { crit, high, medium, normal }*

Default: *normal*

The priority of traffic for a specified address.

4. Purging The MAC Address Table

This command removes all dynamic entries from the MAC address table. The only negative impact of this operation is that it causes flooding while addresses are relearned.

The “Purge MAC Address Table” action (*purge-mac-table*) is accessible from the *mac-tables* menu.

Figure 27.8. Purge MAC Address menu

Figure 27.9. Purge MAC Address Table form

To purge MAC address tables, click the *Perform* button on the Purge MAC Address Table form. This form appears on the same screen as the *purge-mac-table* menu.

28. Spanning Tree

ROX™ provides the latest in IEEE standard Spanning Tree functionality, including:

- Industry standard support of Rapid Spanning Tree (802.1D-2004), which features a compatibility mode with legacy STP (802.1D-1998)
- Industry standard support of Multiple Spanning Trees (802.1Q-2005), which is interoperable with both RSTP and legacy STP.
- RuggedCom RSTP feature enhancements (eRSTP™)
- Superior performance - RSTP will recognize a link failure and put an alternate port into forwarding within milliseconds.
- RSTP may be enabled on a per-port basis.
- Ports may be configured as edge ports, which allow rapid transitioning to the forwarding state for non-STP hosts.
- Path costs may be hard-configured or determined by port speed negotiation, in either the STP or RSTP style.
- Full bridge and port status displays provide a rich set of tools for performance monitoring and debugging.

Historically, a device implementing STP on its ports has been referred to as a bridge. RuggedCom uses the terms "bridge" and "switch" synonymously.

- SNMP-manageable including newRoot and topologyChange traps.

28.1. RSTP Operation

The 802.1D Spanning Tree Protocol (STP) was developed to enable the construction of robust networks that incorporate redundancy while pruning the active topology of the network to prevent loops. While STP is effective, it requires that frame transfer halt after a link outage until all bridges in the network are guaranteed to be aware of the new topology. Using the values recommended by 802.1D, this period lasts 30 seconds.

The Rapid Spanning Tree Protocol (RSTP, IEEE 802.1w) was a further evolution of the 802.1D Spanning Tree Protocol. It replaced the settling period with an active handshake between bridges that guarantees the rapid propagation of topology information throughout the network. RSTP also offers a number of other significant innovations, including:

- Topology changes in RSTP can originate from and be acted upon by any designated bridges, leading to more rapid propagation of address information, unlike topology changes in STP, which must be passed to the root bridge before they can be propagated to the network.
- RSTP explicitly recognizes two blocking roles - Alternate and Backup Port - which are included in computations of when to learn and forward. STP, however, recognizes only one state - Blocking - for ports that should not forward.
- RSTP bridges generate their own configuration messages, even if they fail to receive any from the root bridge. This leads to quicker failure detection. STP, by contrast, must relay configuration messages received on the root port out its designated ports. If an STP bridge fails to receive a message from its neighbor, it cannot be sure where along the path to the root a failure occurred.
- RSTP offers edge port recognition, allowing ports at the edge of the network to forward frames immediately after activation, while at the same time protecting them against loops.

While providing much better performance than STP, IEEE 802.1w RSTP still required up to several seconds to restore network connectivity when a topology change occurred.

A revised and highly optimized RSTP version was defined in the IEEE standard 802.1D-2004 edition. IEEE 802.1D-2004 RSTP reduces network recovery times to just milliseconds and optimizes RSTP operation for various scenarios.

ROX™ supports IEEE 802.1D-2004 RSTP.

28.1.1. RSTP States and Roles

RSTP bridges have roles to play, either root or designated. One bridge - the Root Bridge - is the logical center of the network. All other bridges in the network are Designated bridges.

RSTP also assigns each port of the bridge a state and a role. The RSTP state describes what is happening at the port in relation to address learning and frame forwarding. The RSTP role basically describes whether the port is facing the center or the edges of the network and whether it can currently be used.

State

There are three RSTP states: Discarding, Learning and Forwarding.

The discarding state is entered when the port is first put into service. The port does not learn addresses in this state and does not participate in frame transfer. The port looks for RSTP traffic in order to determine its role in the network. When it is determined that the port will play an active part in the network, the state will change to learning.

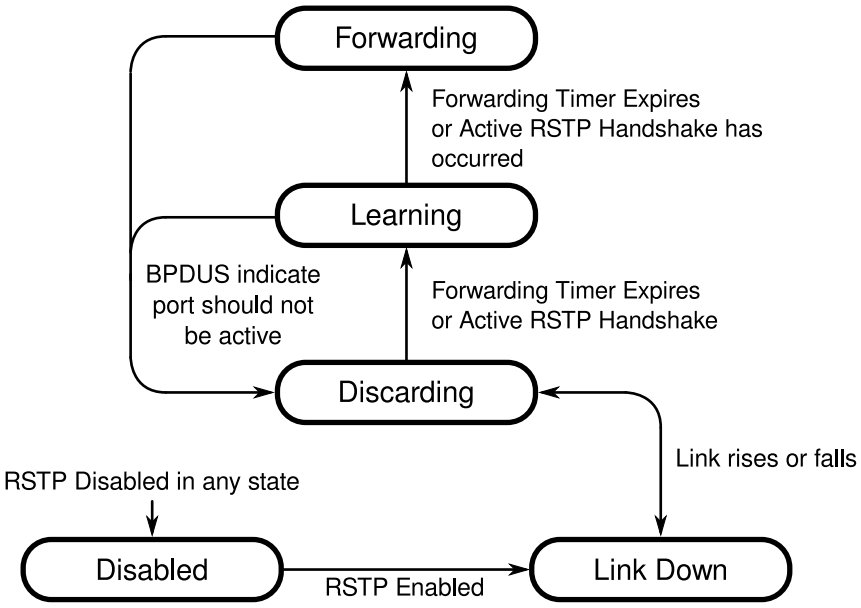

Figure 28.1. Bridge and Port States

The learning state is entered when the port is preparing to play an active part in the network. The port learns addresses in this state but does not participate in frame transfer. In a network of RSTP bridges, the time spent in this state is usually quite short. RSTP bridges operating in STP compatibility mode will spend six to 40 seconds in this state.

After “learning,” the bridge will place the port in the forwarding state. The port both learns addresses and participates in frame transfer while in this state.

ROX™ introduces two more states - Disabled and Link Down. Introduced purely for purposes of management, these states may be considered subclasses of the RSTP Discarding state. The Disabled state refers to links for which RSTP has been disabled. The Link Down state refers to links for which RSTP is enabled but are currently down.

Role

There are four RSTP port roles: Root, Designated, Alternate and Backup.

If the bridge is not the root bridge, it must have a single Root Port. The Root Port is the “best” (i.e. quickest) way to send traffic to the root bridge.

A port is marked as Designated if it is the best port to serve the LAN segment it is connected to. All bridges on the same LAN segment listen to each others’ messages and agree on which bridge is the Designated Bridge. The ports of other bridges on the segment must become either Root, Alternate or Backup ports.

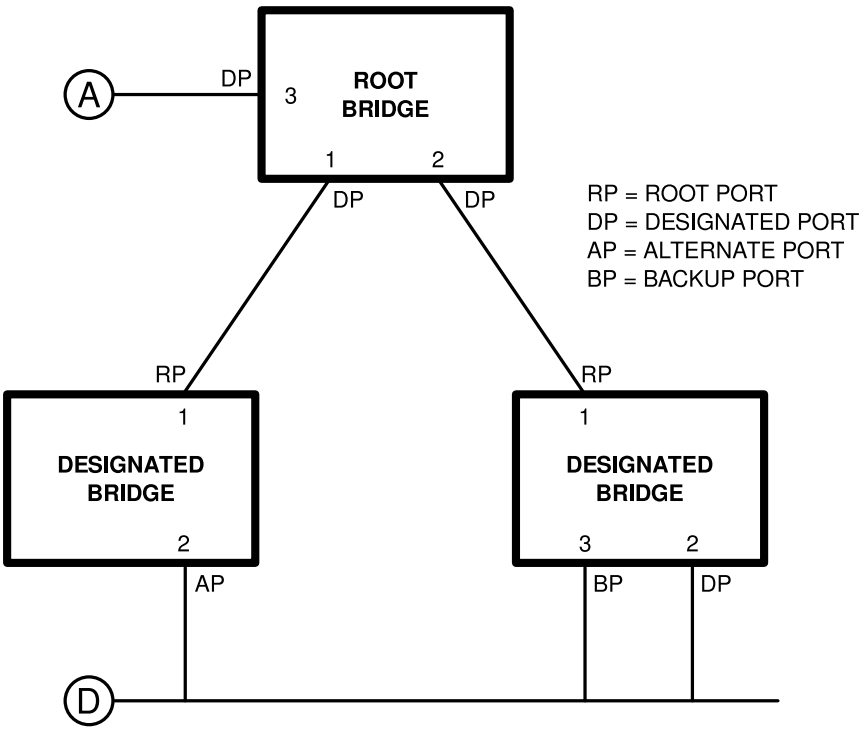

Figure 28.2. Bridge and Port Roles

A port is alternate when it receives a better message from another bridge on the LAN segment it is connected to. The message that an Alternate Port receives is better than the port itself would generate, but not good enough to convince it to become the Root Port. The port becomes the alternate to the current Root Port and will become the new Root Port should the current Root Port fail. The Alternate Port does not participate in the network.

A port is a Backup Port when it receives a better message from the LAN segment it is connected to, originating from another port on the same bridge. The port is a backup for another port on the bridge and will become active if that port fails. The Backup Port does not participate in the network.

28.1.2. Edge Ports

A port may be designated an Edge Port if it is directly connected to an end station. As such, it cannot create bridging loops in the network and can thus directly transition to forwarding, skipping the listening and learning stages.

Edge ports that receive configuration messages immediately lose their Edge Port status and become normal spanning tree ports. A loop created on an improperly connected edge port is thus quickly repaired.

Because an Edge Port services only end stations, topology change messages are not generated when its link toggles.

28.1.3. Point-to-Point and Multipoint Links

RSTP uses a peer-peer protocol called Proposing-Agreeing to ensure transitioning in the event of a link failure. This protocol is point-to-point and breaks down in multipoint situations, i.e. when more than two bridges operate on a shared media link.

If RSTP detects this circumstance (based upon the port's half duplex state after link up) it will switch off Proposing-Agreeing. The port must transition through the learning and forwarding states, spending one forward delay in each state.

There are circumstances in which RSTP will make an incorrect decision about the point-to-point state of the link simply by examining the half-duplex status, namely:

- The port attaches only to a single partner, but through a half-duplex link.
- The port attaches to a shared media hub through a full-duplex link. The shared media link attaches to more than one RSTP enabled bridge.

In such cases, the user may configure the bridge to override the half-duplex determination mechanism and force the link to be treated in the proper fashion.

28.1.4. Path and Port Costs

The STP path cost is the main metric by which root and designated ports are chosen[1]. The path cost for a designated bridge is the sum of the individual port costs of the links between the root bridge and that designated bridge. The port with the lowest path cost is the best route to the root bridge and is chosen as the root port.

How Port Costs Are Generated

Port costs can be generated either as a result of link auto-negotiation or manual configuration.

When the link auto-negotiation method is used, the port cost is derived from the speed of the link. This method is useful when a well-connected network has been established. It can be used when the designer is not too concerned with the resultant topology as long as connectivity is assured.

Manual configuration is useful when the exact topology of the network must be predictable under all circumstances. The path cost can be used to establish the topology of the network exactly as the designer intends.

[1] In actuality the primary determinant for root port selection is the root bridge ID. Bridge ID is important mainly at network startup when the bridge with the lowest ID is elected as the root bridge. After startup (when all bridges agree on the root bridge's ID) the path cost is used to select root ports. If the path costs of candidates for the root port are the same, the ID of the peer bridge is used to select the port. Finally, if candidate root ports have the same path cost and peer bridge ID, the port ID of the peer bridge is used to select the root port. In all cases the lower ID, path cost or port ID is selected as the best.

STP vs. RSTP Costs

The IEEE 802.1D-1998 specification limits port costs to values of 1 to 65536. It recommends that a path cost corresponding to the 1x109 / link speed be used. Designed at a time when 9600 bps links were state of the art, this method breaks down in modern use, as the method cannot represent a link speed higher than a gigabit per second.

In order to remedy this problem in future applications the IEEE 802.1w specification limits port costs to values of 1 to 200000, with a path cost corresponding to the 2x1012 / link speed.

RuggedCom bridges support interoperability with legacy STP bridges by selecting the style to use. In practice it makes no difference which style is used as long as it is applied consistently across the network, or if costs are manually assigned.

28.1.5. Bridge Diameter

The bridge diameter is the maximum number of bridges between any two possible points of attachment of end stations to the network.

The bridge diameter reflects the realization that topology information requires time to propagate hop by hop through a network. If configuration messages take too long to propagate end to end through the network, the result will be an unstable network.

There is a relationship between the bridge diameter and the maximum age parameter[2]. To achieve extended ring sizes, RuggedCom eRSTP™ uses an age increment of ¼ of a second. The value of the maximum bridge diameter is thus four times the configured maximum age parameter.

Raise the value of the maximum age parameter if implementing very large bridged networks or rings.

28.2. MSTP Operation

The Multiple Spanning Tree (MST) algorithm and protocol provide greater control and flexibility than RSTP and legacy STP. MSTP (Multiple Spanning Tree Protocol) is an extension of RSTP, whereby multiple spanning trees may be maintained on the same bridged network. Data traffic is allocated to one or another of several spanning trees by mapping one or more VLANs onto the network.

The sophistication and utility of the Multiple Spanning Tree implementation on a given bridged network is proportional to the amount of planning and design invested in configuring MSTP.

If MSTP is activated on some or all of the bridges in a network with no additional configuration, the result will be a fully and simply connected network, but at best, the result will be the same as a network using only RSTP. Taking full advantage of the features offered by MSTP requires a potentially large number of configuration variables to be derived from an analysis of data traffic on the bridged network, and from requirements for load sharing, redundancy, and path optimization. Once these parameters have all been derived, it is also critical that they are consistently applied and managed across all bridges in an MST region.

[2]The RSTP algorithm is as follows. STP configuration messages contain “age” information. Messages transmitted by the root bridge have an age of 0. As each subsequent designated bridge transmits the configuration message it must increase the age by at least 1 second. When the age exceeds the value of the maximum age parameter the next bridge to receive the message immediately discards it.

28.2.1. MST Regions and Interoperability

In addition to supporting multiple spanning trees in a network of MSTP-capable bridges, MSTP is capable of interoperating with bridges that support only RSTP or legacy STP, without requiring any special configuration.

An MST region may be defined as the set of interconnected bridges whose MST Region Identification is identical. The interface between MSTP bridges and non-MSTP bridges, or between MSTP bridges with different MST Region Identification information, becomes part of an MST Region boundary.

Bridges outside an MST region will see the entire region as though it were a single (R)STP bridge; the internal detail of the MST region is hidden from the rest of the bridged network. In support of this, MSTP maintains separate 'hop counters' for spanning tree information exchanged at the MST region boundary versus that propagated inside the region. For information received at the MST region boundary, the (R)STP Message Age is incremented only once. Inside the region, a separate Remaining Hop Count is maintained, one for each spanning tree instance. The external Message Age parameter is referred to the (R)STP Maximum Age Time, whereas the internal Remaining Hop Counts are compared to an MST region-wide Maximum Hops parameter.

MSTI

An MSTI (Multiple Spanning Tree Instance) is one of sixteen independent spanning tree instances that may be defined in an MST region (not including the IST – see below). An MSTI is created by mapping a set of VLANs (in ROX™, via the VLAN configuration) to a given MSTI ID. The same mapping must be configured on all bridges that are intended to be part of the MSTI. Moreover, all VLAN to MSTI mappings must be identical for all bridges in an MST region.

ROX™ supports 16 MSTIs in addition to the IST.

Each MSTI has a topology that is independent of every other. Data traffic originating from the same source and bound to the same destination but on different VLANs on different MSTIs may therefore travel a different path across the network.

IST

An MST region always defines an IST (Internal Spanning Tree). The IST spans the entire MST region, and carries all data traffic that is not specifically allocated (by VLAN) to a specific MSTI. The IST is always computed and is defined to be MSTI zero.

The IST is also the extension inside the MST region of the CIST (see below), which spans the entire bridged network, inside and outside of the MST region and all other RSTP and STP bridges, as well as any other MST regions.

CST

The CST (Common Spanning Tree) spans the entire bridged network, including MST regions and any connected STP or RSTP bridges. An MST region is seen by the CST as an individual bridge, with a single cost associated with its traversal.

CIST

The CIST (Common and Internal Spanning Tree) is the union of the CST and the ISTs in all MST regions. The CIST therefore spans the entire bridged network, reaching into each MST region via the latter's IST to reach every bridge on the network.

28.2.2. MSTP Bridge and Port Roles

28.2.2.1. Bridge Roles:

CIST Root

The CIST Root is the elected root bridge of the CIST (Common and Internal Spanning Tree), which spans all connected STP and RSTP bridges and MSTP regions.

CIST Regional Root

The root bridge of the IST within an MST region. The CIST Regional Root is the bridge within an MST region with the lowest cost path to the CIST Root. Note that the CIST Regional Root will be at the boundary of an MST region. Note also that it is possible for the CIST Regional Root to be the CIST Root.

MSTI Regional Root

The root bridge for an MSTI within an MST region. A root bridge is independently elected for each MSTI in an MST region.

28.2.2.2. Port Roles:

Each port on an MST bridge may have more than one role depending on the number and topology of spanning tree instances defined on the port.

CIST Port Roles

- The Root Port provides the minimum cost path from the bridge to the CIST Root via the CIST Regional Root. If the bridge itself happens to be the CIST Regional Root, the Root Port is also the Master Port for all MSTIs (see below), and provides the minimum cost path to a CIST Root located outside the region.
- A Designated Port provides the minimum cost path from an attached LAN, via the bridge to the CIST Regional Root.
- Alternate and Backup Ports have the same sense that they do in RSTP, described in Section 28.1.1, “RSTP States and Roles”, under “Roles”, but relative to the CIST Regional Root.

MSTI Port Roles

For each MSTI on a bridge:

- The Root Port provides the minimum cost path from the bridge to the MSTI Regional Root, if the bridge itself is not the MSTI Regional Root.
- A Designated Port provides the minimum cost path from an attached LAN, via the bridge to the MSTI Regional Root.
- Alternate and Backup Ports have the same sense that they do in RSTP, described in Section 28.1.1, “RSTP States and Roles”, under “Roles”, but relative to the MSTI Regional Root.

The Master Port, which is unique in an MST region, is the CIST Root Port of the CIST Regional Root, and provides the minimum cost path to the CIST Root for all MSTIs.

Boundary Ports

A Boundary Port is a port on a bridge in an MST region that connects to either of: 1) a bridge belonging to a different MST region, or 2) a bridge supporting only RSTP or legacy STP. A Boundary Port blocks or forwards all VLANs from all MSTIs and the CIST alike. A Boundary Port may be:

- The CIST Root Port of the CIST Regional Root (and therefore also the MSTI Master Port).

- A CIST Designated Port, CIST Alternate / Backup Port, or Disabled. At the MST region boundary, the MSTI Port Role is the same as the CIST Port Role.

A Boundary Port connected to an STP bridge will send only STP BPDUs. One connected to an RSTP bridge need not refrain from sending MSTP BPDUs. This is made possible by the fact that the MSTP carries the CIST Regional Root Identifier in the field that RSTP parses as the Designated Bridge Identifier.

28.2.3. Benefits of MSTP

Despite the fact that MSTP is configured by default to arrive automatically at a spanning tree solution for each configured MSTI, advantages may be gained from influencing the topology of MSTIs in an MST region. The fact that the Bridge Priority and each port cost are configurable per MSTI (see Section 28.4.4, “Port MSTI Parameters”) makes it possible to control the topology of each MSTI within a region.

Load Balancing

MST can be used to balance data traffic load among (sets of) VLANs, enabling more complete utilization of a multiply interconnected bridged network.

A bridged network controlled by a single spanning tree will block redundant links by design, in order to avoid harmful loops. Using MSTP, however, any given link may have a different blocking state for each spanning tree instance (MSTI), as maintained by MSTP. Any given link, therefore, might be in blocking state for some VLANs and in forwarding state for other VLANs, depending on the mapping of VLANs to MSTIs.

It is possible to control the spanning tree solution for each MSTI, especially the set of active links for each tree, by manipulating, per MSTI, the bridge priority and the port costs of links in the network. If traffic is allocated judiciously to multiple VLANs, redundant interconnections in a bridged network which, using a single spanning tree, would have gone unused, can now be made to carry traffic.

Isolation of Spanning Tree Reconfiguration

A link failure in an MST region that does not affect the roles of Boundary ports will not cause the CST to be reconfigured, nor will the change affect other MST regions. This is due to the fact that MSTP information does not propagate past a region boundary.

MSTP versus PVST

An advantage of MSTP over the Cisco Systems Inc. proprietary PVST protocol is the ability to map multiple VLANs onto a single MSTI. Since each spanning tree requires processing and memory, the expense of keeping track of an increasing number of VLANs increases much more rapidly for PVST than for MSTP.

Compatibility with STP and RSTP

No special configuration is required for the bridges of an MST region to connect fully and simply to non-MST bridges on the same bridged network. Careful planning and configuration is, however, recommended in order to arrive at an optimal network.

28.2.4. Implementing MSTP on a Bridged Network

It is recommended that the configuration of MSTP on a network proceed in the sequence outlined below. Naturally, it is also recommended that network analysis and planning inform the steps of configuring the VLAN and MSTP parameters in particular.

Begin with a set of MSTP-capable Ethernet bridges, and MSTP disabled. For each bridge in the network:

1. Configure and enable RSTP (see Section 28.4.1, “Spanning Tree Parameters” and Section 28.4.2, “Port RSTP Parameters”). Note that the Max Hops parameter in the Bridge RSTP Parameters menu is the maximum hop count for MSTP.
2. Create the VLANs that will be mapped to MSTIs (see the sections on VLAN Configuration).
3. Map VLANs to MSTIs (via the VLAN Configuration menus). Note that MSTP need not be enabled in order to map a VLAN to an MSTI. Note also that this mapping must be identical for each bridge that is to belong to the MST region.
4. Configure a Region Identifier and Revision Level. Note that these two items must be identical for each bridge in the MST region .
5. Configure Bridge Priority per MSTI .
6. Configure Port Cost and Priority per port and per MSTI (see Section 28.4.4, “Port MSTI Parameters”).
7. Enable MSTP (see Section 28.4.1, “Spanning Tree Parameters”).

Static VLANs must be used in an MSTP configuration. GVRP is not supported in this case.

28.3. RSTP Applications

28.3.1. RSTP in Structured Wiring Configurations

RSTP allows you to construct structured wiring systems in which connectivity is maintained in the event of link failures. For example, a single link failure of any of links A through N in Figure 28.3, “Example of a Structured Wiring Configuration”, would leave all the ports of bridges 555 through 888 connected to the network.

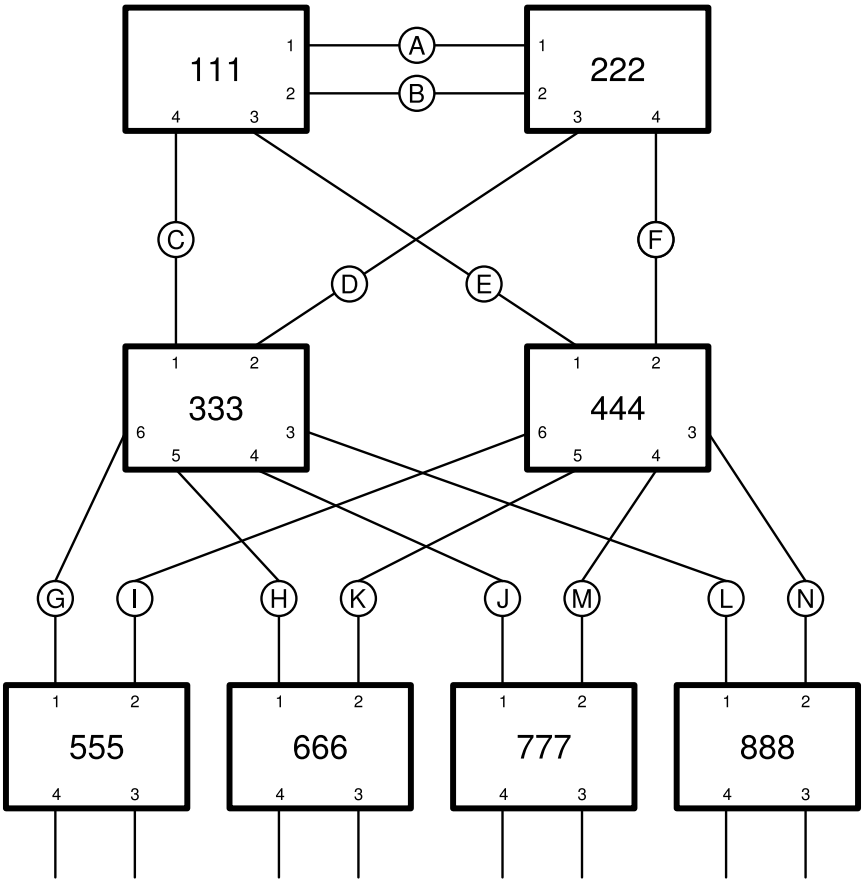

Figure 28.3. Example of a Structured Wiring Configuration

Procedure 28.1. Design Considerations for RSTP in Structured Wiring Configurations

1. **Select the design parameters for the network.**

   What are the requirements for robustness and network fail-over/recovery times? Are there special requirements for diverse routing to a central host computer? Are there any special port redundancy requirements?

2. **Identify required legacy support.**

   Are STP bridges used in the network? These bridges do not support rapid transitioning to forwarding. If these bridges are present, can they be re-deployed closer to the network edge?

3. **Identify edge ports and ports with half-duplex/shared media restrictions.**

   Ports that connect to host computers, IEDs and controllers may be set to edge ports in order to guarantee rapid transitioning to forwarding as well as to reduce the number of topology change notifications in the network. Ports with half-duplex/shared media restrictions require special attention in order to guarantee that they do not cause extended fail-over/recovery times.

4. **Choose the root bridge and backup root bridge carefully.**

   The root bridge should be selected to be at the concentration point of network traffic. Locate the backup root bridge adjacent to the root bridge. One strategy that may be used is to tune the bridge priority to establish the root bridge and then tune each bridge's priority to correspond to its distance from the root bridge.

5. **Identify desired steady state topology.**

   Identify the desired steady state topology taking into account link speeds, offered traffic and QOS. Examine of the effects of breaking selected links, taking into account network loading and the quality of alternate links.

6. **Decide upon port cost calculation strategy.**

   Select whether fixed or auto-negotiated costs should be used? Select whether the STP or RSTP cost style should be used.

7. **Calculate and configure priorities and costs.**

8. **Implement the network and test under load.**

28.3.2. RSTP in Ring Backbone Configurations

RSTP may be used in ring backbone configurations where rapid recovery from link failure is required. In normal operation, RSTP will block traffic on one of the links, for example as indicated by the double bars through link H in Figure 28.4, “Example of a Ring Backbone Configuration”. In the event of a failure on link D, bridge 444 will unblock link H. Bridge 333 will communicate with the network through link F.

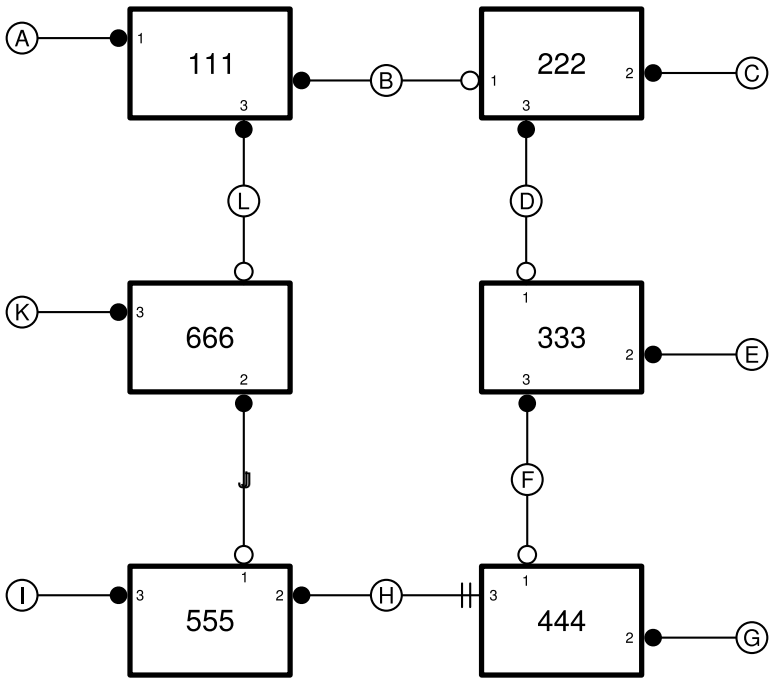

Figure 28.4. Example of a Ring Backbone Configuration

Procedure 28.2. Design Considerations for RSTP in Ring Backbone Configurations

1. **Select the design parameters for the network.**

   What are the requirements for robustness and network fail-over/recovery times? Typically, ring backbones are chosen to provide cost effective but robust network designs.

2. **Identify required legacy support and ports with half-duplex/shared media restrictions.**

   These bridges should not be used if network fail-over/recovery times are to be minimized.

3. **Identify edge ports**

   Ports that connect to host computers, IEDs and controllers may be set to edge ports in order to guarantee rapid transitioning to forwarding as well as to reduce the number of topology change notifications in the network.

4. **Choose the root bridge.**

   The root bridge can be selected to equalize either the number of bridges, number of stations or amount of traffic on either of its legs. It is important to realize that the ring will always be broken in one spot and that traffic always flows through the root.

5. **Assign bridge priorities to the ring.**

   The strategy that should be used is to assign each bridge's priority to correspond to its distance from the root bridge. If the root bridge is assigned the lowest priority of 0, the bridges on either side should use a priority of 4096 and the next bridges 8192 and so on. As there are 16 levels of bridge priority available, this method provides for up to 31 bridges in the ring.

6. **Implement the network and test under load.**

28.3.3. RSTP Port Redundancy

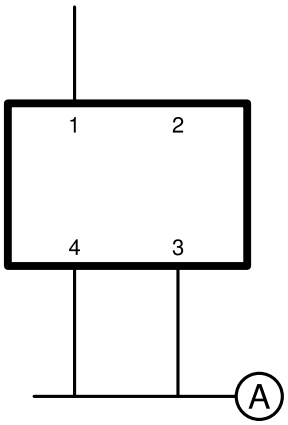

Figure 28.5. Port Redundancy

In cases where port redundancy is essential, RSTP allows more than one bridge port to service a LAN. For example, if port 3 is designated to carry the network traffic of LAN A, port 4 will block. Should an interface failure occur on port 3, port 4 would assume control of the LAN.

28.4. Spanning Tree Configuration

The Spanning Tree menu is accessible from the main menu under *switch*. The path to this menu is *switch/spanning-tree*. The RSTP Common Instanc, Spanning Tree Parameter form and eRSTP form appear on the same screen as the *Spanning Tree* menu.

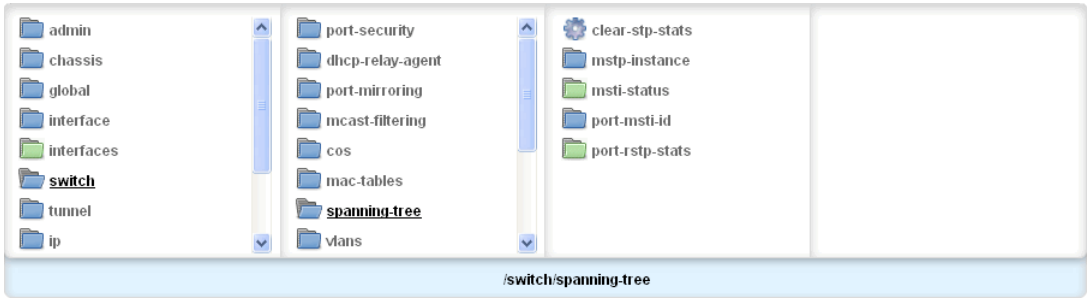

Figure 28.6. Spanning Tree menu

28.4.1. Spanning Tree Parameters

The Spanning Tree parameter form at the top-level Spanning Tree menu configures parameters applicable to RSTP and MSTP over the whole bridge.

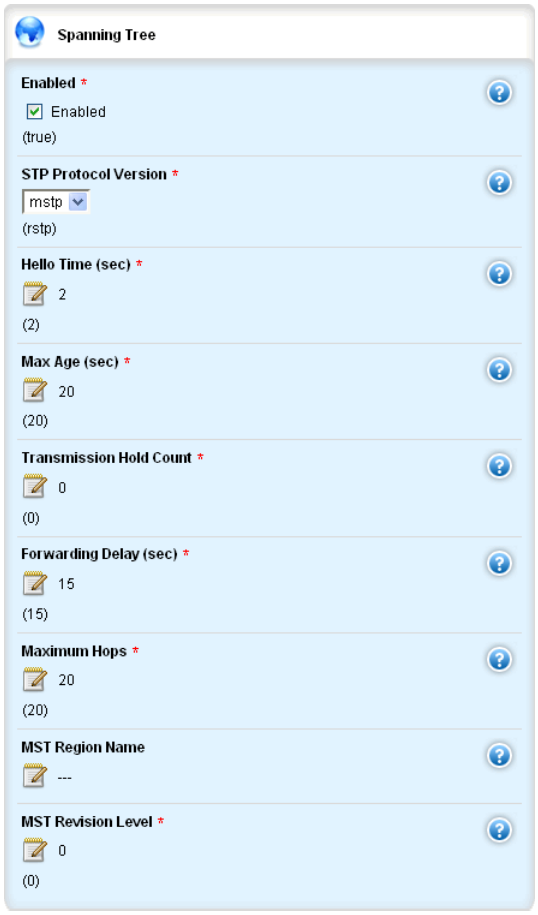

Figure 28.7. Spanning Tree Parameter form

Enabled

Synopsis: *boolean*

Default: *true*

Enables STP/RSTP/MSTP for the bridge globally. Note that STP/RSTP/MSTP is enabled on a port when it is enabled globally and along with enabling per port setting.

STP Protocol Version

Synopsis: *string - one of the following keywords { mstp, rstp, stp }*

Default: *rstp*

The version of the Spanning Tree Protocol to support, either only STP or Rapid STP or Multiple STP

Hello Time (sec)

Synopsis: *unsigned integer*

Default: *2*

The time between configuration messages issued by the root bridge. Shorter hello times result in faster detection of topology changes at the expense of moderate increases in STP traffic. (Relationship : maxAgeTime >= 2 * (helloTime + 1.0 seconds))

Max Age (sec)

Synopsis: *unsigned integer*

Default: *20*

The time for which a configuration message remains valid after being issued by the root bridge. Configure this parameter with care when many tiers of bridges exist, or slow speed links (such as those used in WANs) are part of the network. (Relationship : maxAgeTime >= 2 * (helloTime + 1.0 seconds))

Transmission Hold Count

Synopsis: *unsigned integer*

Default:

The maximum number of configuration messages on each port that may be sent in a special event (such as recovering from a failure or bringing up a new link). After the maximum number of messages is reached, RSTP will be limited to 1 message per second. Larger values allow the network to recover from failed links more quickly. If RSTP is being used in a ring architecture, the transmit count should be larger than the number of switches in the ring. If unset, then the value is interpreted as unlimited (default).

Forwarding Delay (sec)

Synopsis: *unsigned integer*

Default: *15*

The amount of time a bridge spends learning MAC addresses on a rising port before beginning to forward traffic. Lower values allow the port to reach the forwarding state more quickly, but at the expense of flooding unlearned addresses to all ports.

Maximum Hops

Synopsis: *unsigned integer*

Default: *20*

This parameter is only relevant for MSTP; ignore it otherwise. This parameter specifies the maximum possible bridge diameter inside an MST region. MSTP BPDUs propagating inside an MST region carry a time-to-live parameter decremented by every switch that propagates the BPDU. If the maximum number of hops inside the region exceeds the configured maximum, BPDUs may be discarded due to their time-to-live information.

MST Region Name

Synopsis: *A string*

This is used to configure a new region from a subset of switches in a current region, while maintaining the same region name.

MST Revision Level

Synopsis: *unsigned integer*

Default:

Variable length text string. You must configure an identical region name on all switches you want to be in the same MST region.

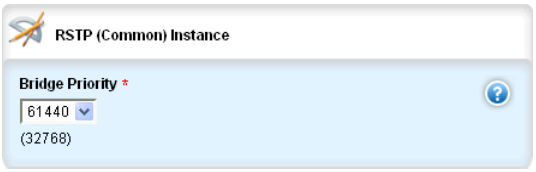

Figure 28.8. RSTP Common Instance form

Bridge Priority

Synopsis: *string - one of the following keywords { 61440, 57344, 53248, 49152, 45960, 40960, 36864, 32768, 28672, 24576, 20480, 16384, 12288, 8192, 4096, 0 }*

Default: *32768*

The priority assigned to the RSTP / Common Bridge Instance

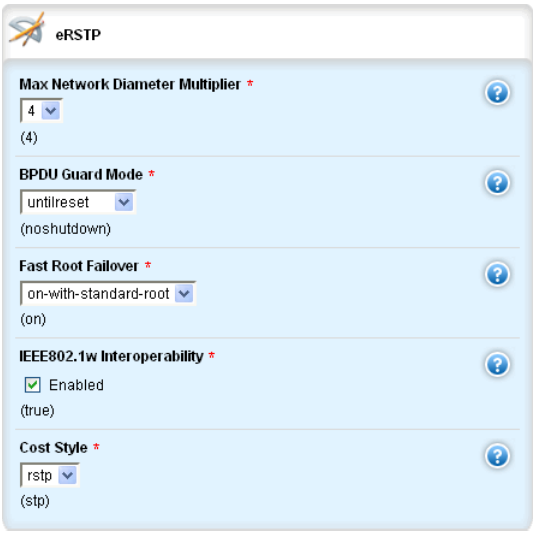

Figure 28.9. eRSTP form

Set Enhanced RSTP Parameters using the eRSTP form.

Max Network Diameter Multiplier

Synopsis: *string - one of the following keywords { 4, 1 }*

Default: *4*

The Max Network Diameter as a muliplier of the MaxAgeTime value

BPDU Guard Mode

Synopsis: *string - one of the following keywords { untilreset, noshutdown, specify }*

Default: *noshutdown*

The RSTP standard does not address network security. RSTP must process every received BPDU and take an appropriate action. This opens a way for an attacker to influence RSTP topology by injecting RSTP BPDUs into the network. BPDU Guard is a feature that protects the network from BPDUs received by a port where RSTP-capable devices are not expected to be attached. If a BPDU is received by a port for which the 'Edge' parameter is set to 'TRUE' or RSTP is disabled, the port will be shut down for the time period specified by this parameter.

- NO SHUTDOWN : BPDU Guard is disabled.

- UNTIL RESET : The port will remain shut down until the port reset command is issued by the user.
- SPECIFY : a timeout period is specified for the port.

BPDU Timeout

Synopsis: *unsigned integer*

The RSTP standard does not address network security. RSTP must process every received BPDU and take an appropriate action. This opens a way for an attacker to influence RSTP topology by injecting RSTP BPDUs into the network. BPDU Guard protects the network from BPDUs received by a port where RSTP-capable devices are not expected to be attached. If a BPDU is received by a port for which the 'Edge' parameter is set to 'TRUE' or RSTP is disabled, the port will be shut down for the time period specified by this parameter.

- DON'T SHUTDOWN : BPDU Guard is disabled.
- UNTIL RESET : The port will remain shut down until the port reset command is issued by the user.

Fast Root Failover

Synopsis: *string - one of the following keywords { on-with-standard-root, off, on }*

Default: *on*

In mesh network topologies, the standard RSTP algorithm does not guarantee deterministic network recovery time in the case of a root switch failure. Such a recovery time is hard to calculate and it can be different (and may be relatively long) for any given mesh topology. This configuration parameter enables RuggedCom's enhancement to RSTP which detects a failure of the root switch and performs some extra RSTP processing steps, significantly reducing the network recovery time and making it deterministic.

This feature is only available in RSTP mode. In MSTP mode, the configuration parameter is ignored. In a single ring topology, this feature is not needed and should be disabled to avoid longer network recovery times due to extra RSTP processing. The Fast Root Failover algorithm must be supported by all switches in the network, including the root, to guarantee optimal performance. However, it is not uncommon to assign the root role to a switch from a vendor different from the rest of the switches in the network. In other words, it is possible that the root might not suport the Fast Root Failover algorithm. In such a scenario,a relaxed algorithm should be used, which tolerates the lack of support in the root switch.

These are the supported configuration options:

- Off: Fast Root Failover algorithm is disabled and hence a root switch failure; may result in excessive connectivity recovery time.
- On: Fast Root Failover is enabled and the most robust algorithm is used, which requires the appropriate support in the root switch.
- On with standard root: Fast Root Failover is enabled but a relaxed algorithm is used, allowing the use of a standard switch in the root role.

IEEE802.1w Interoperability

Synopsis: *boolean*

Default: *true*

IEEE802.1w Interoperability

Cost Style

Synopsis: *string - one of the following keywords { rstp, stp }*

Default: *stp*

The style of link costs to employ. STP uses 16-bit path costs based upon 1x10E9/link speed (4 for 1Gbps, 19 for 100 Mbps and 100 for 10 Mbps) whereas RSTP uses 32-bit costs based upon 2x10E13/link speed (20,000 for 1Gbps, 200,000 for 100 Mbps and 2,000,000 for 10 Mbps). Note that RSTP link costs are used only when the bridge version support is set to allow RSTP and the port does not migrate to STP.

28.4.2. Port RSTP Parameters

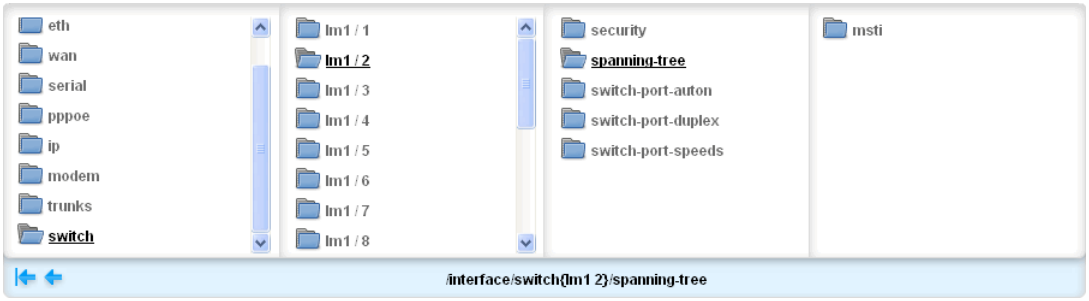

Figure 28.10. Interface/switch/{line module}/spanning-tree submenu

This submenu is accessible from the main menu under *interface/switch/{line module}/spanning-tree*.

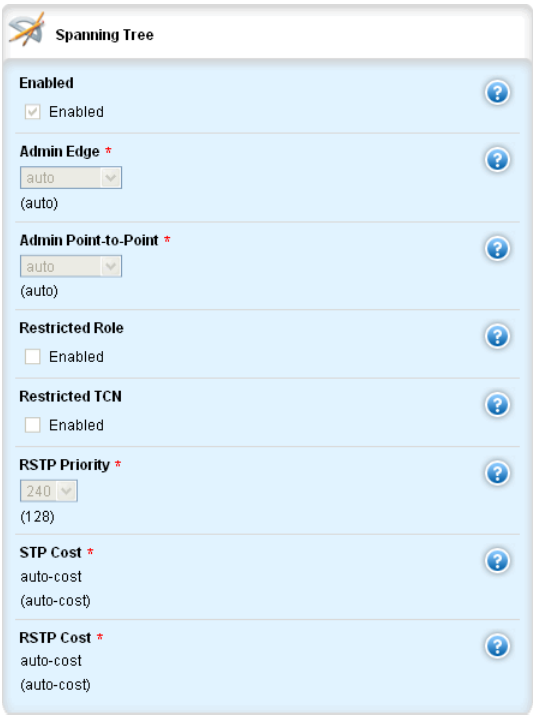

Figure 28.11. Port RSTP Parameter form

The Port RSTP Parameter form appears on the same screen as the *interface/switch/{line module}/spanning-tree* submenu.

Enabled

Synopsis: *boolean*

Default: *true*

When the box is checked, the Spanning Tree Protocol is enabled on the interface. Enabling STP activates the STP or RSTP protocol for this interface per the configuration in the STP Configuration menu.

Admin Edge

Synopsis: *string - one of the following keywords { auto, forceFalse, forceTrue }*

Default: *auto*

Edge ports are ports that do not participate in the Spanning Tree, but still send configuration messages. Edge ports transition directly to frame forwarding without any listening and learning delays. The MAC tables of Edge ports do not need to be flushed when topology changes occur in the STP network. Unlike an STP-disabled port, accidentally connecting an edge port to another port in the spanning tree will result in a detectable loop. The 'Edgeness' of the port will be switched off and the standard RSTP rules will apply (until the next link outage).

Admin Point-to-Point

Synopsis: *string - one of the following keywords { auto, forceFalse, forceTrue }*

Default: *auto*

RSTP uses a peer to peer protocol that provides for rapid transitioning on point-to-point links. This protocol is automatically turned off in situations where multiple STP bridges communicate over a shared (non point-to-point) LAN. The bridge will automatically take point-to-point to be true when the link is found to be operating in full-duplex mode. The point-to-point parameter allows this behavior or overrides it, forcing point-to-point to be true or false. Force the parameter true when the port operates a point-to-point link but cannot run the link in full-duplex mode. Force the parameter false when the port operates the link full duplex, but is still not point to point (e.g. a full-duplex link to an unmanaged bridge that concentrates two other STP bridges)

Restricted Role

If enabled, causes the port not to be selected as the root port for the CIST or any MSTI, even it has the best spanning tree priority vector. This parameter should be FALSE by default.

Restricted TCN

If TRUE, causes the port not to propagate received topology change notifications and topology changes to other ports. This parameter should be FALSE by default. If set it can cause temporary loss of connectivity after changes in a spanning tree's active topology as a result of persistent incorrectly learned station location information.

RSTP Priority

Synopsis: *string - one of the following keywords { 240, 224, 208, 192, 176, 160, 144, 128, 112, 96, 64, 32, 16, 0 }*

Default: *128*

The STP port priority. Ports of the same cost that attach to a common LAN will select the port to be used based upon the port priority.

STP Cost

Synopsis: *string - the keyword { auto-cost }*

Synopsis: *unsigned integer*

Default: *auto-cost*

The cost to use in cost calculations, when the cost style parameter is set to STP in the bridge RSTP parameter configuration. Setting the cost manually provides the ability to preferentially select specific ports to carrytraffic over others. Leave this field set to 'auto' to use the standard STP port costs as negotiated (four for 1Gbps, 19 for 100 Mbps links and 100 for 10 Mbps links). For MSTP, this parameter applies to both external and internal path costs.

RSTP Cost

Synopsis: *string - the keyword { auto-cost }*

Synopsis: *unsigned integer*

Default: *auto-cost*

The cost to use in cost calculations, when the cost style parameter is set to RSTP in the bridge RSTP parameter configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to 'auto' to use the standard RSTP

port costs as negotiated (20,000 for 1Gbps, 200,000 for 100 Mbps links and 2,000,000 for 10 Mbps links). For MSTP, this parameter applies to both external and internal path costs.

28.4.3. Bridge MSTI Parameters

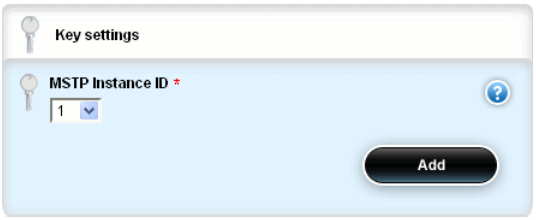

Figure 28.12. Key Settings form

To configure parameters using the Key Settings form and MSTP Instance form, navigate to *switch/spanning-tree/mstp-instance*.

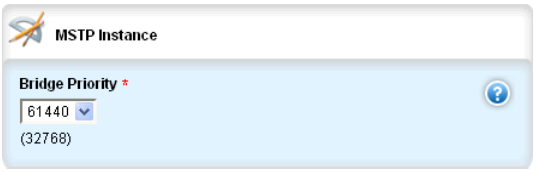

Figure 28.13. MSTP Instance form

MSTP Instance ID

Synopsis: *integer*

MSTP Instance ID

Bridge Priority

Synopsis: *string - one of the following keywords { 61440, 57344, 53248, 49152, 45960, 40960, 36864, 32768, 28672, 24576, 20480, 16384, 12288, 8192, 4096, 0 }*

Default: *32768*

Bridge Priority provides a way to control the topology of the STP connected network. The desired root and designated bridges can be configured for a particular topology. The bridge with the lowest priority will become the root. In the event of a failure of the root bridge, the bridge with the next lowest priority will then become the root. Designated bridges that (for redundancy purposes) service a common LAN also use priority to determine which bridge is active. In this way, careful selection of bridge priorities can establish the path of traffic flows in normal and abnormal conditions.

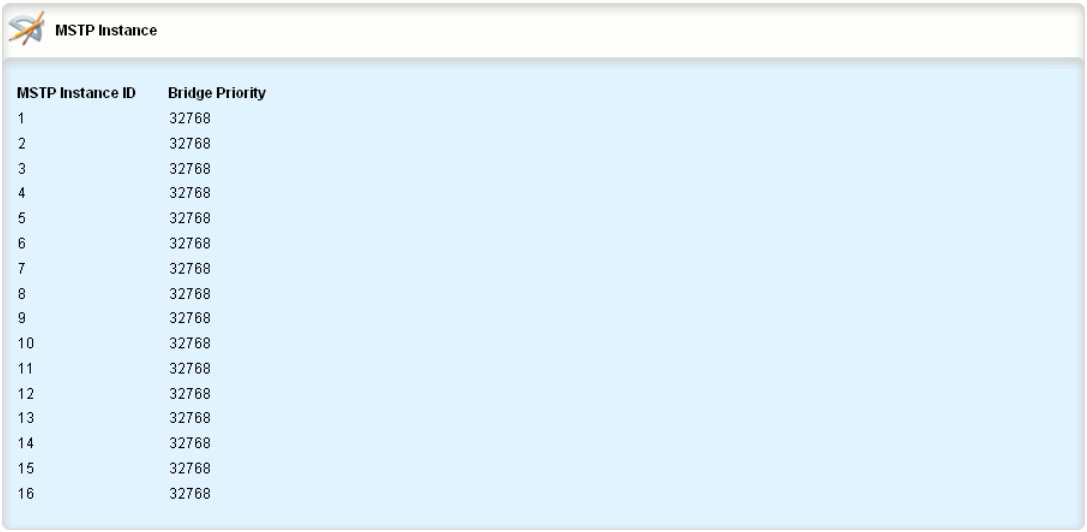
MSTP Instance

MSTP Instance ID	Bridge Priority
1	32768
2	32768
3	32768
4	32768
5	32768
6	32768
7	32768
8	32768
9	32768
10	32768
11	32768
12	32768
13	32768
14	32768
15	32768
16	32768

Figure 28.14. MSTP Instance table

After data has been configured, the MSTP Instance table will be displayed at *switch/spanning-tree/mstp-instance*.

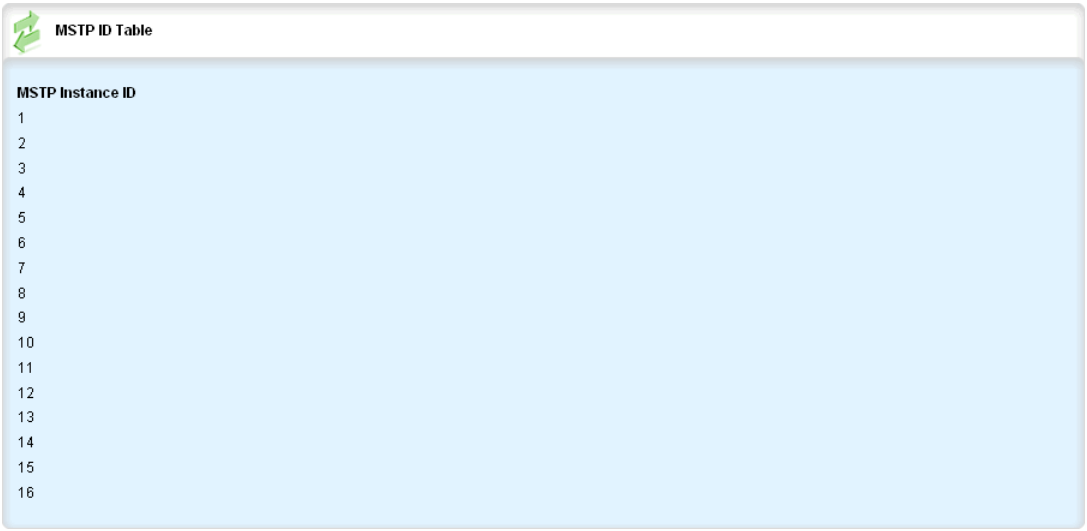
MSTP ID Table

MSTP Instance ID
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16

Figure 28.15. MSTP ID table

To display the MSTP ID table, navigate to *switch/spanning-tree/port-msti-id*.

MSTP Instance ID

Synopsis: *integer*

The MSTP Instance ID.

28.4.4. Port MSTI Parameters

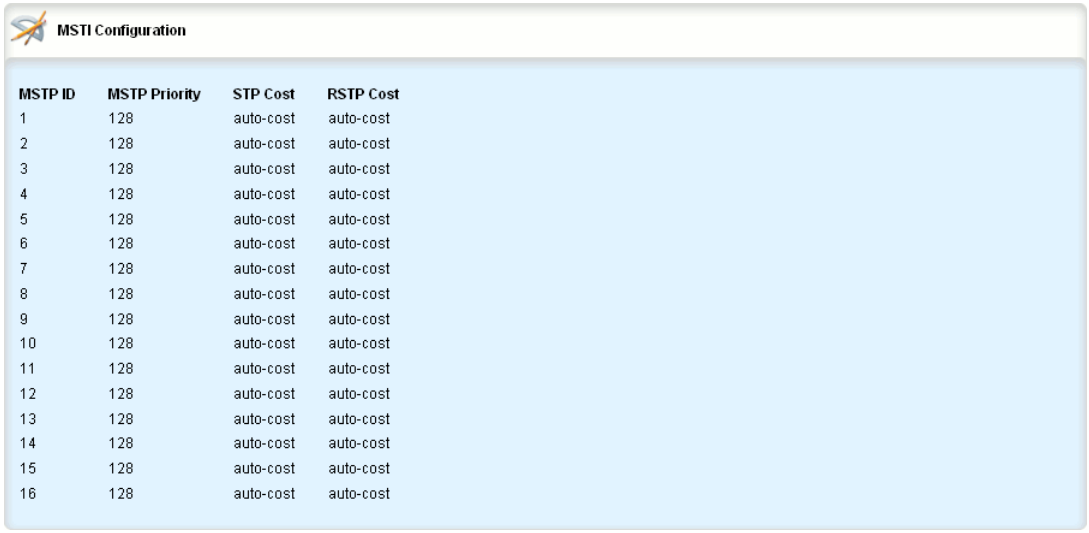
MSTI Configuration
MSTP ID MSTP Priority STP Cost RSTP Cost
1 128 auto-cost auto-cost
2 128 auto-cost auto-cost
3 128 auto-cost auto-cost
4 128 auto-cost auto-cost
5 128 auto-cost auto-cost
6 128 auto-cost auto-cost
7 128 auto-cost auto-cost
8 128 auto-cost auto-cost
9 128 auto-cost auto-cost
10 128 auto-cost auto-cost
11 128 auto-cost auto-cost
12 128 auto-cost auto-cost
13 128 auto-cost auto-cost
14 128 auto-cost auto-cost
15 128 auto-cost auto-cost
16 128 auto-cost auto-cost

Figure 28.16. MSTI Configuration table

To display the MSTI Configuration table, navigate to *interface/switch/{line module}/spanning-tree/msti*.

Figure 28.17. MSTI Configuration form

To display the MSTI Configuration form, navigate to *interface/switch/{line module}/spanning-tree/msti/{number}*.

MSTP ID

Synopsis: *integer*

MSTP Instance Identifier

MSTP Priority

Synopsis: *string - one of the following keywords { 240, 224, 208, 192, 176, 160, 144, 128, 112, 96, 64, 32, 16, 0 }*

Default: *128*

The STP port priority. Ports of the same cost that attach to a common LAN will select the port to be used based upon the port priority.

STP Cost

Synopsis: *string - the keyword { auto-cost }*

Synopsis: *unsigned integer*

Default: *auto-cost*

The cost to use in cost calculations, when the cost style parameter is set to STP in the bridge RSTP parameter configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to 'auto' to use the standard STP port costs as negotiated (four for 1Gbps, 19 for 100 Mbps links and 100 for 10 Mbps links). For MSTP, this parameter applies to both external and internal path costs.

RSTP Cost

Synopsis: *string - the keyword { auto-cost }*

Synopsis: *unsigned integer*

Default: *auto-cost*

The cost to use in cost calculations, when the cost style parameter is set to RSTP in the bridge RSTP parameter configuration. Setting the cost manually provides the ability to preferentially select specific ports to carry traffic over others. Leave this field set to 'auto' to use the standard RSTP port costs as negotiated (20,000 for 1Gbps, 200,000 for 100 Mbps links and 2,000,000 for 10 Mbps links). For MSTP, this parameter applies to both external and internal path costs.

28.5. Spanning Tree Statistics

28.5.1. Bridge RSTP Statistics

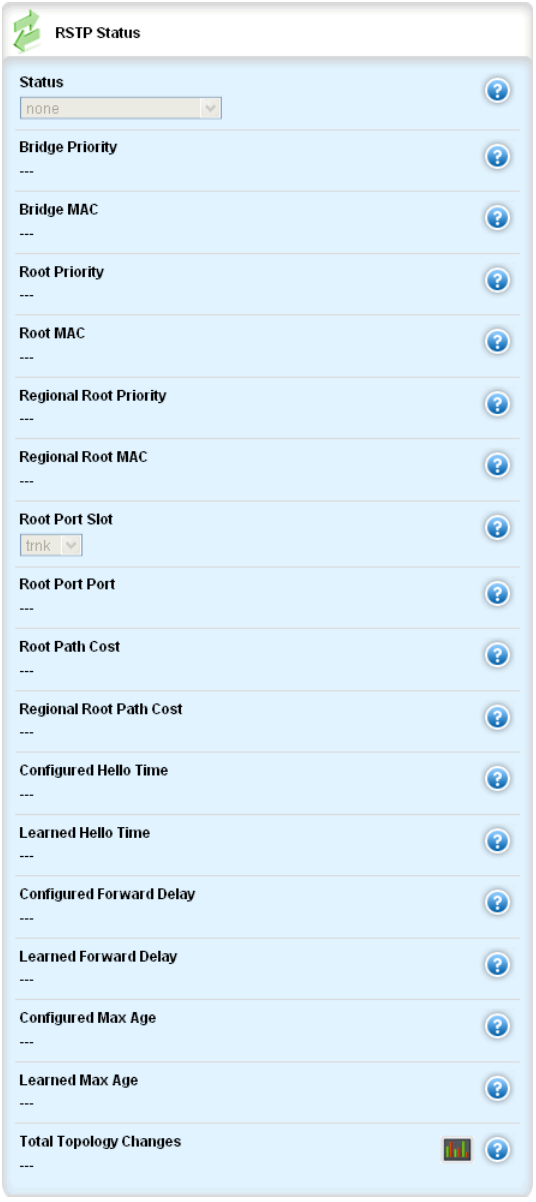

Figure 28.18. RSTP Status form

To display this form, navigate to *switch/spanning-tree*.

Status

Synopsis: *string - one of the following keywords { none, rootBridge, notDesignatedForAnyLAN, designatedBridge }*

The spanning tree status of the bridge. The status may be root or designated. This field may show text saying 'not designated for any LAN' if the bridge is not the designated bridge for any of its ports.

Bridge Priority

Synopsis: *integer*

The bridge identifier of this bridge.

Bridge MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The bridge identifier of this bridge.

Root Priority

Synopsis: *integer*

Ports to which the multicast group traffic is forwarded.

Root MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

Ports to which the multicast group traffic is forwarded.

Regional Root Priority

Synopsis: *integer*

The bridge identifier of the IST regional root bridge for the MST region this device belongs to.

Regional Root MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The bridge identifier of the IST regional root bridge for the MST region this device belongs to.

Root Port Slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - the keyword { trnk }*

If the bridge is designated, this is the slot containing the port that provides connectivity towards the root bridge of the network.

Root Port Port

Synopsis: *integer*

If the bridge is designated, this is the port of the slot that provides connectivity towards the root bridge of the network.

Root Path Cost

Synopsis: *unsigned integer*

The total cost of the path to the root bridge, composed of the sum of the costs of each link in the path. If custom costs have not been configured. 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports will contribute 100. For the CIST instance of MSTP, this is an external root path cost, which is the cost of the path from the IST root (i.e. regional root) bridge to the CST root (i.e. network "global" root) bridge.

Regional Root Path Cost

Synopsis: *unsigned integer*

For the CIST instance of MSTP, this is the cost of the path to the IST root (i.e. regional root) bridge

Configured Hello Time

Synopsis: *integer*

The configured Hello time from the Bridge RSTP Parameters menu.

Learned Hello Time

Synopsis: *integer*

The actual Hello time provided by the root bridge as learned in configuration messages. This time is used in designated bridges.

Configured Forward Delay

Synopsis: *integer*

The configured Forward Delay time from the Bridge RSTP Parameters menu.

Learned Forward Delay

Synopsis: *integer*

The actual Forward Delay time provided by the root bridge as learned in configuration messages. This time is used in designated bridges.

Configured Max Age

Synopsis: *integer*

The configured Maximum Age time from the Bridge RSTP Parameters menu.

Learned Max Age

Synopsis: *integer*

The actual Maximum Age time provided by the root bridge as learned in configuration messages. This time is used in designated bridges.

Total Topology Changes

Synopsis: *unsigned integer*

A count of topology changes in the network, as detected on this bridge through link failures or as signaled from other bridges. Excessively high or rapidly increasing counts signal network problems.

28.5.2. Port RSTP Statistics

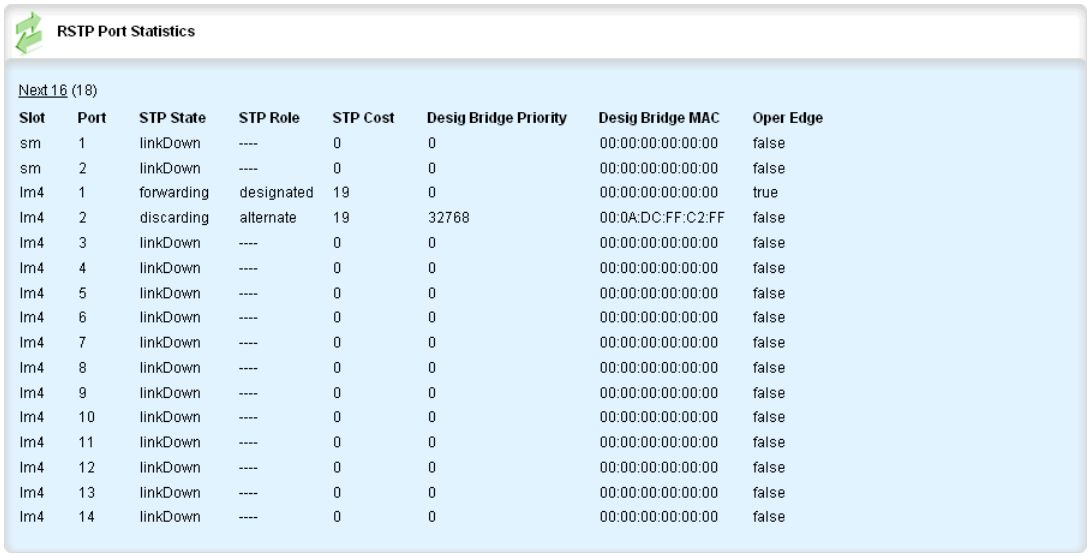
RSTP Port Statistics

Next 16 (18)

Slot	Port	STP State	STP Role	STP Cost	Desig Bridge Priority	Desig Bridge MAC	Oper Edge
sm	1	linkDown	----	0	0	00:00:00:00:00:00	false
sm	2	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	1	forwarding	designated	19	0	00:00:00:00:00:00	true
lm4	2	discarding	alternate	19	32768	00:0A:DC:FF:C2:FF	false
lm4	3	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	4	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	5	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	6	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	7	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	8	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	9	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	10	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	11	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	12	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	13	linkDown	----	0	0	00:00:00:00:00:00	false
lm4	14	linkDown	----	0	0	00:00:00:00:00:00	false

Figure 28.19. RSTP Port Statistics table

To display the RSTP Port Statistics table, navigate to *switch/spanning-tree/port-rstp-stats*.

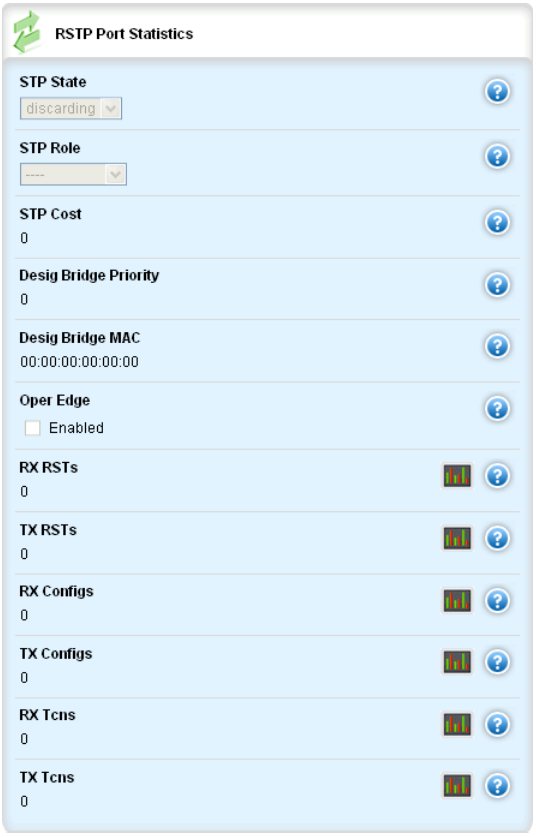

Figure 28.20. RSTP Port Statistics form

To display these forms, navigate to *switch/spanning-tree/port-rstp-stats/{line module}*.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - the keyword { trnk }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the module.

STP State

Synopsis: *string - one of the following keywords { discarding, linkDown, forwarding, learning, listening, blocking, disabled }*

Describes the status of this interface in the spanning tree:

- Disabled : STP is disabled on this port.
- Link Down : STP is enabled on this port but the link is down.
- Discarding : The link is not used in the STP topology but is standing by.
- Learning : The port is learning MAC addresses in order to prevent flooding when it begins forwarding traffic.
- Forwarding : The port is forwarding traffic.

STP Role

Synopsis: *string - one of the following keywords { ----, master, backup, alternate, designated, root }*

The role of this port in the spanning tree:

- Designated : The port is designated for (i.e. carries traffic towards the root for) the LAN it is connected to.
- Root : The single port on the bridge, which provides connectivity towards the root bridge.
- Backup : The port is attached to a LAN that is serviced by another port on the bridge. It is not used but is standing by.
- Alternate : The port is attached to a bridge that provides connectivity to the root bridge. It is not used but is standing by.
- Master : Only exists in MSTP. The port is an MST region boundary port and the single port on the bridge, which provides connectivity for the Multiple Spanning Tree Instance towards the Common Spanning Tree root bridge (i.e. this port is the root port for the Common Spanning Tree Instance).

STP Cost

Synopsis: *unsigned integer*

The cost offered by this port. If the Bridge RSTP Parameters Cost Style is set to STP, 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports contribute 100. If the Cost Style is set to RSTP, 1Gbps will contribute 20,000, 100 Mbps ports will contribute a cost of 200,000 and 10 Mbps ports contribute a cost of 2,000,000. Note that even if the Cost style is set to RSTP, a port that migrates to STP will have its cost limited to a maximum of 65535.

Desig Bridge Priority

Synopsis: *integer*

Provided on the root ports of the designated bridges, the Bridge Identifier of the bridge this port is connected to.

Desig Bridge MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

Provided on the root ports of designated bridges, the Bridge Identifier of the bridge this port is connected to.

Oper Edge

Synopsis: *boolean*

Whether or not the port is operating as an edge port.

RX RSTs

Synopsis: *unsigned integer*

The count of RSTP configuration messages received on this port.

TX RSTs

Synopsis: *unsigned integer*

The count of RSTP configuration messages transmitted on this port.

RX Configs

Synopsis: *unsigned integer*

The count of STP configuration messages received on this port.

TX Configs

Synopsis: *unsigned integer*

The count of STP configuration messages transmitted on this port.

RX Tcns

Synopsis: *unsigned integer*

The count of configuration change notification messages received on this port. Excessively high or rapidly increasing counts signal network problems.

TX Tcns

Synopsis: *unsigned integer*

The count of configuration messages transmitted from this port.

28.5.3. MSTI Status

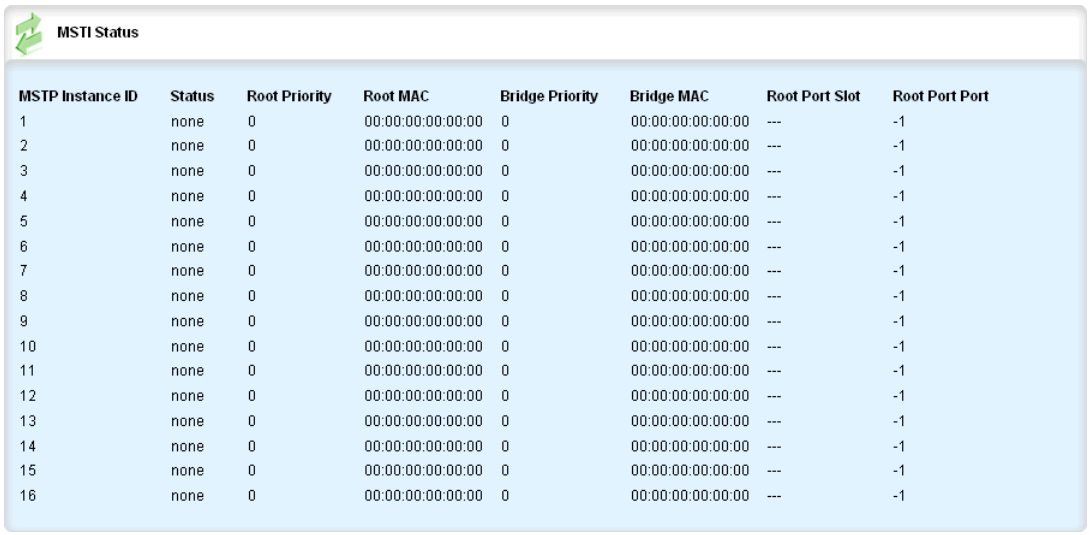

MSTP Instance ID	Status	Root Priority	Root MAC	Bridge Priority	Bridge MAC	Root Port Slot	Root Port Port
1	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
2	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
3	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
4	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
5	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
6	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
7	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
8	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
9	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
10	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
11	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
12	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
13	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
14	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
15	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1
16	none	0	00:00:00:00:00:00	0	00:00:00:00:00:00	---	-1

Figure 28.21. MSTI Status table

To display this table, navigate to *switch/spanning-tree/msti-status*.

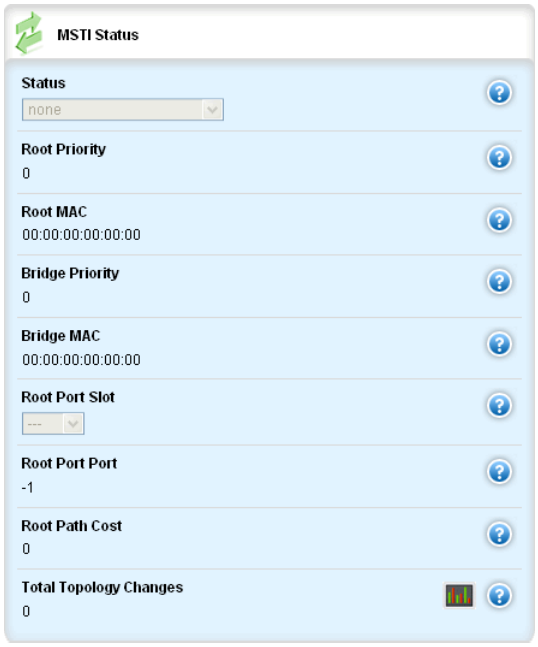

Figure 28.22. MSTI Status form

To display these forms, navigate to *switch/spanning-tree/msti-status/{number}*.

MSTP Instance ID

Synopsis: *integer*

The bridge identifier of this bridge.

status

Synopsis: *string - one of the following keywords { none, rootBridge, notDesignatedForAnyLAN, designatedBridge }*

The spanning tree status of the bridge. The status may be root or designated. This field may show text saying 'not designated for any LAN' if the bridge is not the designated bridge for any of its ports.

Root Priority

Synopsis: *integer*

Bridge Identifier of the root bridge.

Root MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

Bridge Identifier of the root bridge.

Bridge Priority

Synopsis: *integer*

Bridge Identifier of this bridge.

Bridge MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

Bridge Identifier of this bridge.

Root Port Slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - the keyword { trnk }*

If the bridge is designated, this is the slot containing the port that provides connectivity towards the root bridge of the network.

Root Port Port

Synopsis: *integer*

If the bridge is designated, this is the port of the slot that provides connectivity towards the root bridge of the network.

Root Path Cost

Synopsis: *unsigned integer*

The total cost of the path to the root bridge composed of the sum of the costs of each link in the path. If custom costs have not been configured. 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports will contribute 100. For the CIST instance of MSTP, this is an external root path cost, which is the cost of the path from the IST root (i.e. regional root) bridge to the CST root (i.e. network "global" root) bridge.

Total Topology Changes

Synopsis: *unsigned integer*

A count of topology changes in the network, as detected on this bridge through link failures or as signaled from other bridges. Excessively high or rapidly increasing counts signal network problems.

28.5.4. Port MSTP Statistics

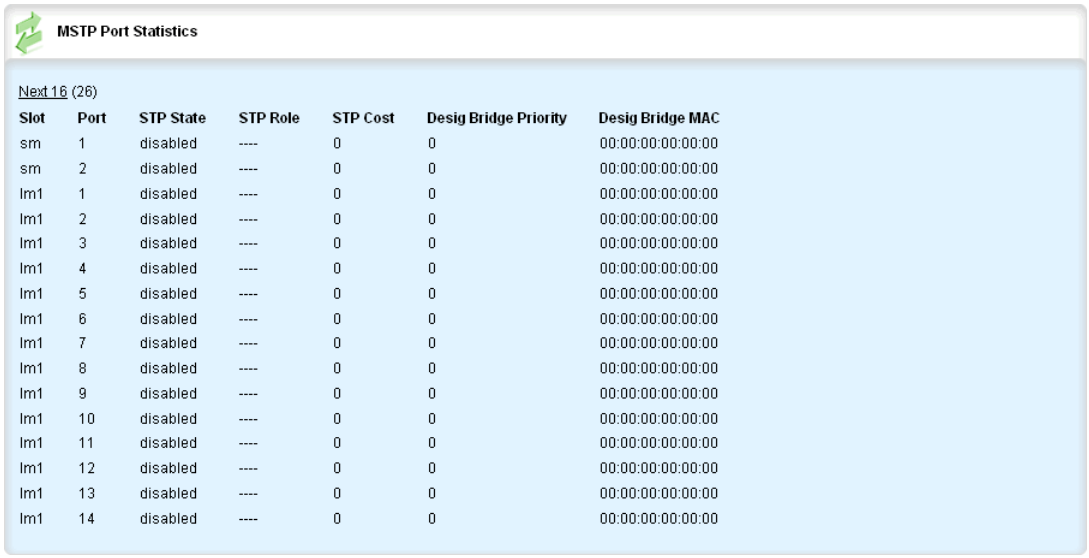

Slot	Port	STP State	STP Role	STP Cost	Desig Bridge Priority	Desig Bridge MAC
sm	1	disabled	----	0	0	00:00:00:00:00:00
sm	2	disabled	----	0	0	00:00:00:00:00:00
lm1	1	disabled	----	0	0	00:00:00:00:00:00
lm1	2	disabled	----	0	0	00:00:00:00:00:00
lm1	3	disabled	----	0	0	00:00:00:00:00:00
lm1	4	disabled	----	0	0	00:00:00:00:00:00
lm1	5	disabled	----	0	0	00:00:00:00:00:00
lm1	6	disabled	----	0	0	00:00:00:00:00:00
lm1	7	disabled	----	0	0	00:00:00:00:00:00
lm1	8	disabled	----	0	0	00:00:00:00:00:00
lm1	9	disabled	----	0	0	00:00:00:00:00:00
lm1	10	disabled	----	0	0	00:00:00:00:00:00
lm1	11	disabled	----	0	0	00:00:00:00:00:00
lm1	12	disabled	----	0	0	00:00:00:00:00:00
lm1	13	disabled	----	0	0	00:00:00:00:00:00
lm1	14	disabled	----	0	0	00:00:00:00:00:00

Figure 28.23. MSTP Port Statistics table

The path to the MSTP Port Statistics table is *switch/spanning-tree/port-msti-id/{number}/port-msti-stats*.

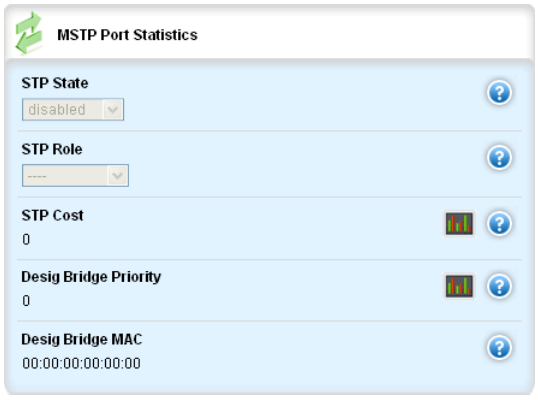

Figure 28.24. MSTP Port Statistics form

The path to MSTP Port Statistics forms is *switch/spanning-tree/port-msti-id/{number}/port-msti-stats/{line module}*.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - the keyword { trnk }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the module.

STP State

Synopsis: *string - one of the following keywords { discarding, linkDown, forwarding, learning, listening, blocking, disabled }*

The status of this interface in the spanning tree:

- Disabled : STP is disabled on this port.
- Link Down : STP is enabled on this port but the link is down.
- Discarding : The link is not used in the STP topology but is standing by.
- Learning : The port is learning MAC addresses in order to prevent flooding when it begins forwarding traffic.
- Forwarding : The port is forwarding traffic.

STP Role

Synopsis: *string - one of the following keywords { ----, master, backup, alternate, designated, root }*

The role of this port in the spanning tree:

- Designated : The port is designated for (i.e. carries traffic towards the root for) the LAN it is connected to.
- Root : The single port on the bridge, which provides connectivity towards the root bridge.
- Backup : The port is attached to a LAN that is serviced by another port on the bridge. It is not used but is standing by.
- Alternate : The port is attached to a bridge that provides connectivity to the root bridge. It is not used but is standing by.
- Master : Only exists in MSTP. The port is an MST region boundary port and the single port on the bridge, which provides connectivity for the Multiple Spanning Tree Instance towards the Common Spanning Tree root bridge (i.e. this port is the root port for the Common Spanning Tree Instance).

STP Cost

Synopsis: *unsigned integer*

The total cost of the path to the root bridge composed of the sum of the costs of each link in the path. If custom costs have not been configured. 1Gbps ports will contribute a cost of four, 100 Mbps ports will contribute 19 and 10 Mbps ports will contribute 100. For the CIST instance of MSTP, this is an external root path cost, which is the cost of the path from the IST root (i.e. regional root) bridge to the CST root (i.e. network "global" root) bridge.

Desig Bridge Priority

Synopsis: *integer*

The bridge identifier of this bridge.

Desig Bridge MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The bridge identifier of this bridge.

28.6. Clearing Spanning Tree Statistics

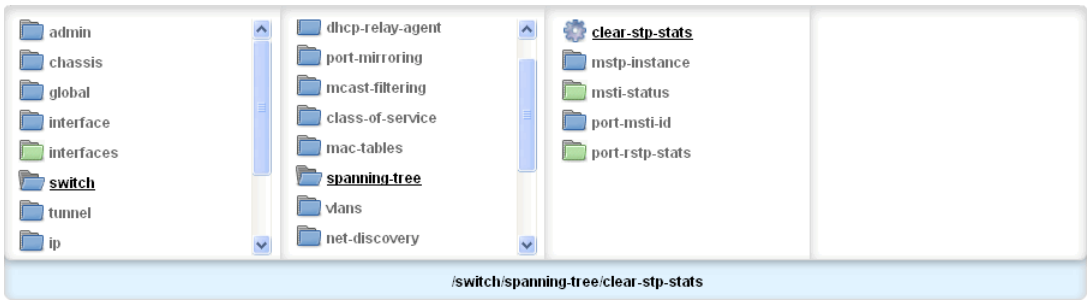

Figure 28.25. Clear-stp-stats Menu Action

The Spanning-Tree Statistics form clears all spanning tree statistics for ethernet ports. This form is accessible from the clear-stp-stats menu action. The path to this menu action is *switch/spanning-tree/clear-stp-stats*. To clear statistics, click the *Perform* button on the Clear Spanning-Tree Statistics form.

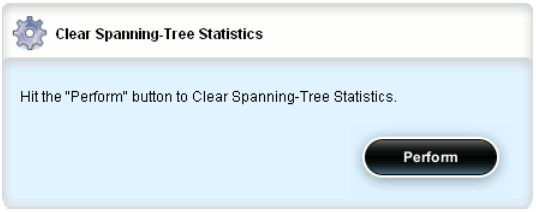

Figure 28.26. Clear Spanning-Tree Statistics form

28.7. Troubleshooting

Problem One

When I connect a new port the network locks up. The port status LEDs are flashing madly.

Occasionally, the network seems to experience a lot of flooding. All the ports seem to experience significant traffic. The problem lasts a few seconds and then goes away.

One of my switches displays a strange behavior where the root port hops back and forth between two switch ports and never settles down.

Is it possible that one of the switches in the network or one of the ports on a switch in the network has STP disabled and accidentally connects to another switch? If this has occurred, then a traffic loop has been formed.

If the problem appears to be transient in nature, it is possible that ports that are part of the spanning tree have been configured as edge ports. After the link layers have come up on edge ports, STP will directly transition them (perhaps improperly) to the forwarding state. If an RSTP configuration message is then received, the port will be returned to blocking. A traffic loop may be formed for the length of time the port was in forwarding.

If one of the switches appears to flip the root from one port to another, the problem may be one of traffic prioritization (See problem five).

Another possible cause of intermittent operation is that of an auto-negotiation mismatch. If one end of the link is fixed to full-duplex mode and the peer auto-negotiates, the auto-negotiating end will fall back to half-duplex operation. At lower traffic, the volumes the link may display few if any errors. As the traffic volume rises, the fixed negotiation side will begin to experience dropped packets while the auto-negotiating side will experience collisions. Ultimately, as traffic loads approach 100%, the link will become entirely unusable. At this point, RSTP will not be able to transmit configuration messages over the link and the spanning tree topology will break down. If an alternate trunk exists, RSTP will activate it in the place of the congested port. Since activation of the alternate port often relieves the congested port of its traffic, the congested port will once again become reliable. RSTP will promptly enter it back into service, beginning the cycle once again. The root port will flip back and forth between two ports on the switch.

Problem Two

My PC/IED/Device is connected to your switch. After I reset the switch, it takes a long time before it comes up.

Is it possible that the RSTP edge setting for this port is set to false? If Edge is set to false, the bridge will make the port go through two forward delay times before the port can send or receive frames. If Edge is set to true, the bridge will transition the port directly to forwarding upon link up.

Another possible explanation is that some links in the network run in half-duplex mode. RSTP uses a peer-to-peer protocol called Proposal-Agreement to ensure transitioning in the event of a link failure. This protocol requires full-duplex operation. When RSTP detects a non-full duplex port, it cannot rely on Proposal-Agreement protocol and must make the port transition the slow (i.e. STP) way. If possible, configure the port for full-duplex operation. Otherwise, configure the port's point-to-point setting to true. Either one will allow the Proposal-Agreement protocol to be used.

Problem Three

When I test your switch by deliberately breaking a link, it takes a long time before I can poll devices past the switch. I thought RSTP was supposed to be fast. What is happening?

Is it possible that some ports participating in the topology have been configured to STP mode or that the port's point-to-point parameter is set to false? STP and multipoint ports converge slowly after failures occur.

Is it possible that the port has migrated to STP? If the port is connected to the LAN segment by shared media and STP bridges are connected to that media, then convergence after link failure will be slow.

Delays on the order of tens or hundreds of milliseconds can result in circumstances where the link broken is the sole link to the root bridge and the secondary root bridge is poorly chosen. The worst of all possible designs occurs when the secondary root bridge is located at the farthest edge of the network from the root. In this case, a configuration message will have to propagate out to the edge and then back in order to reestablish the topology.

Problem Four

My network is composed of a ring of bridges, of which two (connected to each other) are managed and the rest are unmanaged. Why does the RSTP protocol work quickly when I break a link between the managed bridges but not in the unmanaged bridge part of the ring?

A properly operating unmanaged bridge is transparent to STP configuration messages. The managed bridges will exchange configuration messages through the unmanaged bridge part of the ring as if it is non-existent. When a link in the unmanaged part of the ring fails however, the managed bridges will only be able to detect the failure through timing out of hello messages. Full connectivity will require three hello times plus two forwarding times to be restored.

Problem Five

The switch is up and running and working fine. Then I start a certain application and the network becomes unstable. After I stop the application, the network goes back to running normally.

RSTP sends its configuration messages using the highest possible priority level. If CoS is configured to allow traffic flows at the highest priority level and these traffic flows burst continuously to 100% of the line bandwidth, STP may be disrupted. It is therefore advised not to use the highest CoS.

Problem Six

After I bring up a new port, the root moves on to that port, and I don't want it to. The port that I want to become root won't do so.

Is it possible that the port cost is incorrectly programmed or that auto-negotiation derives an undesired value? Inspect the port and path costs with each port active as root.

Problem Seven

My IED/Controller doesn't work with your switch.

Certain low CPU bandwidth controllers have been found to behave less than perfectly when they receive unexpected traffic. Try disabling STP for the port.

If the controller fails around the time of a link outage then there is the remote possibility that frame disordering or duplication may be the cause of the problem. Try setting the root port of the failing controller's bridge to STP.

Problem Eight

My network runs fine with your switch but I occasionally lose polls to my devices.

Inspect network statistics to determine whether the root bridge is receiving TCNs around the time of observed frame loss. It may be possible that you have problems with intermittent links in your network.

Problem Nine

I'm getting a lot of TCNs at the root, where are they coming from?

Examine the RSTP port statistics to determine the port from which the TCNs are arriving. Sign-on to the switch at the other end of the link attached to that port. Repeat this step until the switch generating the TCNs is found (i.e. the switch that is itself not receiving a large number of TCNs). Determine the problem at that switch.

29. Virtual LANs

ROX™ provides the following VLAN features:

- Support for up to 255 VLANs
- Configurable port-native VLAN.
- Port modes of operation tailored to edge devices (such as a PC or IED) and to network switch interconnections.
- A default setting that ensures configuration-free connectivity in certain scenarios.
- Ability to force either tagged or untagged operation on the port-native VLAN.
- GARP VLAN Registration Protocol (GVRP).

29.1. VLAN Operation

29.1.1. VLANs and Tags

A virtual LAN or VLAN is a group of devices on one or more LAN segments that communicate as if they were attached to the same physical LAN segment. VLANs are extremely flexible because they are based on logical instead of physical connections.

When VLANs are introduced, all traffic in the network must belong to one or another VLAN. Traffic on one VLAN cannot pass to another, except through an internetwork router or Layer 3 switch.

A VLAN tag is the identification information that is present in frames in order to support VLAN operation.

29.1.2. Tagged vs. Untagged Frames

Tagged frames are frames with 802.1Q (VLAN) tags that specify a valid VLAN identifier (VID). Untagged frames are frames without tags or frames that carry 802.1p (prioritization) tags only having prioritization information and a VID of 0. Frames with a VID=0 are also called priority-tagged frames.

When a switch receives a tagged frame, it extracts the VID and forwards the frame to other ports in the same VLAN.

29.1.3. Native VLAN

Each port is assigned a native VLAN number, the Port VLAN ID (PVID). When an untagged frame ingresses a port, it is associated with the port's native VLAN.

By default, when the switch transmits a frame on the native VLAN, it sends the frame untagged. The switch can be configured to transmit frames on the native VLAN tagged.

29.1.4. Edge and Trunk Port Types

Each port can be configured as an Edge or Trunk port:

Edge Port

An edge port attaches to a single end device, such as a PC or IED. An edge port carries traffic on a single pre-configured VLAN: the native VLAN.

Trunk Port

Trunk ports are part of the network and carry traffic for all VLANs between switches. Trunk ports are automatically members of all VLANs configured in the switch.

The switch can "pass through" traffic, forwarding frames received on one trunk port out of another trunk port. The trunk ports must be members of all VLANs that the "pass through" traffic is part of, even if none of those VLANs are used on edge ports.

Frames transmitted out of the port on all VLANs other than the port's native VLAN are always sent tagged.

> *Sometimes it may be desirable to manually restrict the traffic on the trunk to a specified group of VLANs; for example, when the trunk connects to a device, such as a Layer 3 router, that supports a subset of the available VLANs. To prevent the trunk port from being a member of the VLAN, include it in the VLAN's Forbidden Ports list.*

Port Type	VLANs Supported	PVID Format	Usage
Edge	1 (Native) Configured	Untagged	*VLAN Unaware networks* – All frames are sent and received without the need for VLAN tags.
		Tagged	*VLAN Aware networks* – VLAN traffic domains are enforced on a single VLAN.
Trunk	All Configured	Tagged or Untagged	*Switch-to-Switch connections* – VLANs must be manually created and administered or can be dynamically learned through GVRP. *Multiple-VLAN end devices* – Implement connections to end devices that support multiple VLANs at the same time.

Table 29.1. Port Types

29.1.5. VLAN Ingress and Egress Rules

Ingress Rules

The VLAN ingress rules are applied to all frames when they are received by the switch:

Frame received This does not depend on ingress port's VLAN configuration parameters	Untagged	Priority Tagged (VID=0)	Tagged (valid VID)
VLAN ID associated with the frame	PVID	PVID	VID in the tag
Frame dropped due to its tagged/untagged format	No	No	No
Frame dropped, if frame associated with VLAN not configured (or learned) in the switch	N/A	N/A	Yes
Frame dropped, if ingress port is not a member of the VLAN the frame is associated with	N/A	N/A	No

Table 29.2. Ingress Rules

Egress Rules

The VLAN egress rules are applied to all frames when they are transmitted by the switch:

Frame sent Egress port type	On egress port's native VLAN	On other VLAN	
		Port is a member of the VLAN	Port is not a member of the VLAN
Edge	According to the egress port's "PVID Format" parameter	N/A (frame is dropped)	
Trunk		Tagged	dropped

Table 29.3. Egress Rules

29.1.6. Forbidden Ports List

Each VLAN can be configured to exclude ports from membership in the VLAN.

29.1.7. VLAN-aware Mode of Operation

The native operation mode for an IEEE 802.1Q compliant switch is VLAN-aware. Even if a specific network architecture does not use VLANs, ROX™ default VLAN settings allow the switch still to

operate in a VLAN-aware mode while providing functionality required for almost any network application. However, the IEEE 802.1Q standard defines a set of rules that must be followed by all VLAN-aware switches:

- Valid VID range is 1 to 4094 (VID=0 and VID=4095 are invalid).
- Each frame ingressing a VLAN-aware switch is associated with a valid VID.
- Each frame egressing a VLAN-aware switch is either untagged or tagged with a valid VID (this means priority-tagged frames with VID=0 are never sent out by a VLAN-aware switch).

Some applications have requirements conflicting with the IEEE 802.1Q native mode of operation. For example, some applications explicitly require priority-tagged frames to be received by end devices).

29.1.8. GVRP (GARP VLAN Registration Protocol)

GVRP is a standard protocol built on GARP (the Generic Attribute Registration Protocol) to automatically distribute VLAN configuration information in a network. Each switch in a network needs only to be configured with VLANs it requires locally; it dynamically learns the rest of the VLANs configured elsewhere in the network via GVRP. A GVRP-aware end station, configured for a particular VLAN ID, can be connected to a trunk on a GVRP-aware switch and automatically become part of the desired VLAN.

When a switch sends GVRP BPDUs out of all GVRP-enabled ports, GVRP BPDUs advertise all the VLANs known to that switch (configured manually or learned dynamically through GVRP) to the rest of the network.

When a GVRP-enabled switch receives a GVRP BPDU advertising a set of VLANs, the receiving port becomes a member of those advertised VLANs and the switch begins advertising those VLANs via all the GVRP-enabled ports (other than the port on which the VLANs were learned).

To improve network security using VLANs, GVRP-enabled ports may be configured to prohibit the learning of any new dynamic VLANs but at the same time be allowed to advertise the VLANs configured on the switch.

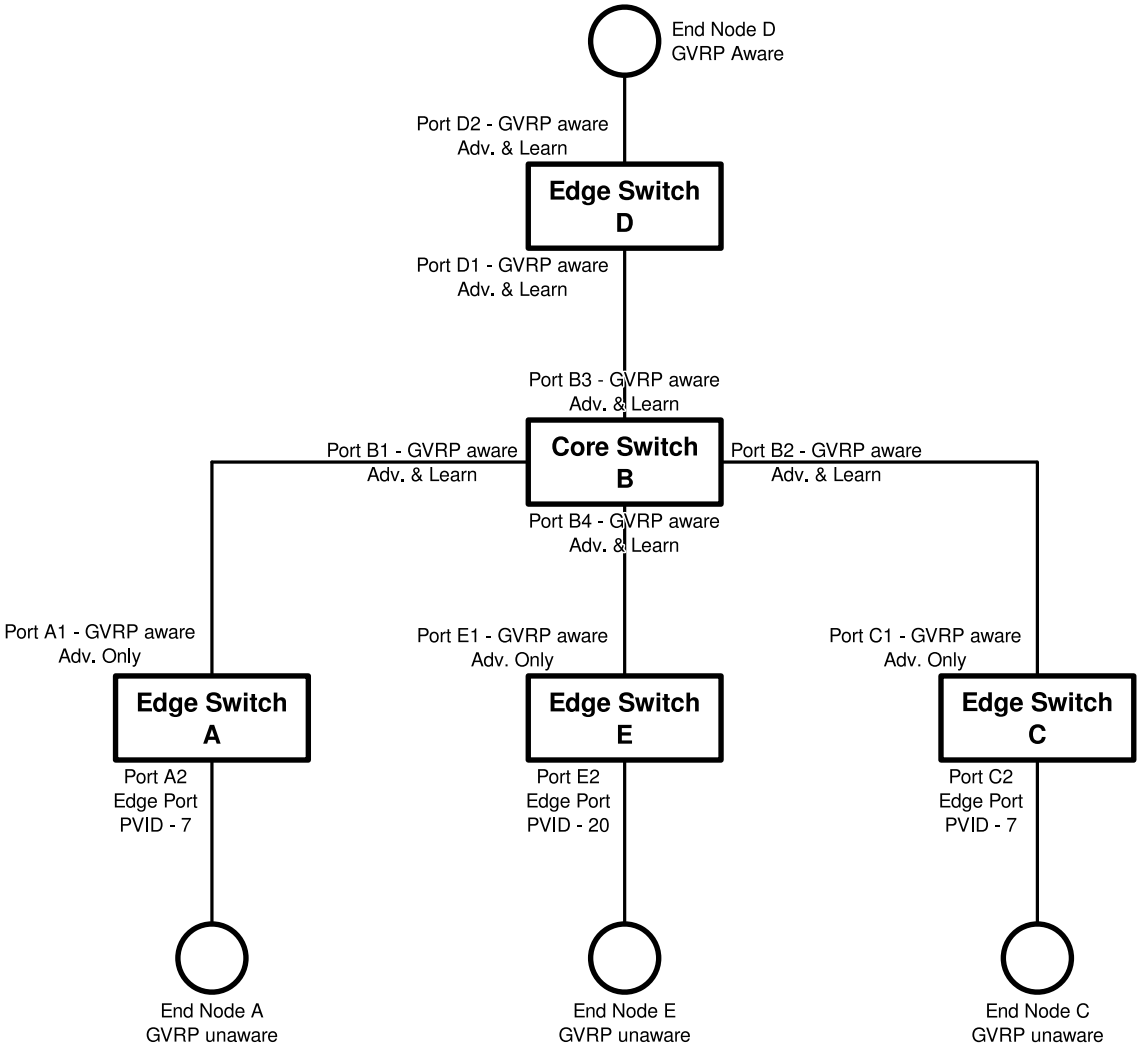

Figure 29.1. Using GVRP

An example of using GVRP:

- Ports A2, and C2 are configured with PVID 7 and port E2 is configured with PVID 20.
- End Node D is GVRP aware and is interested in VLAN 20, hence VLAN 20 is advertised by it towards switch D.
- D2 becomes member of VLAN 20.
- Ports A1 and C1 advertise VID 7 and ports B1 and B2 become members of VLAN 7.
- Ports D1 and B1 advertise VID 20 and ports B3, B4 and D1 become members of VLAN 20.

29.1.9. PVLAN Edge

PVLAN Edge (Protected VLAN Edge port) refers to a feature of the switch whereby multiple VLAN Edge ports on a single device are effectively isolated from one another. All VLAN Edge ports in a switch that are configured as "protected" in this way are prohibited from sending frames to each other, but are still allowed to send frames to other, non-protected, ports within the same VLAN. This protection extends to all traffic on the VLAN: unicast, multicast, or broadcast.

Note that this feature is strictly local to the switch. PVLAN Edge ports are not prevented from communicating with ports off the switch, whether protected (remotely) or not.

29.2. VLAN Applications

29.2.1. Traffic Domain Isolation

VLANs are most often used for their ability to restrict traffic flows between groups of devices.

Unnecessary broadcast traffic can be restricted to the VLAN that requires it. Broadcast storms in one VLAN need not affect users in other VLANs.

Hosts on one VLAN can be prevented from accidentally or deliberately assuming the IP address of a host on another VLAN.

The use of creative bridge filtering and multiple VLANs can carve seemingly unified IP subnets into multiple regions policed by different security/access policies.

Multi-VLAN hosts can assign different traffic types to different VLANs.

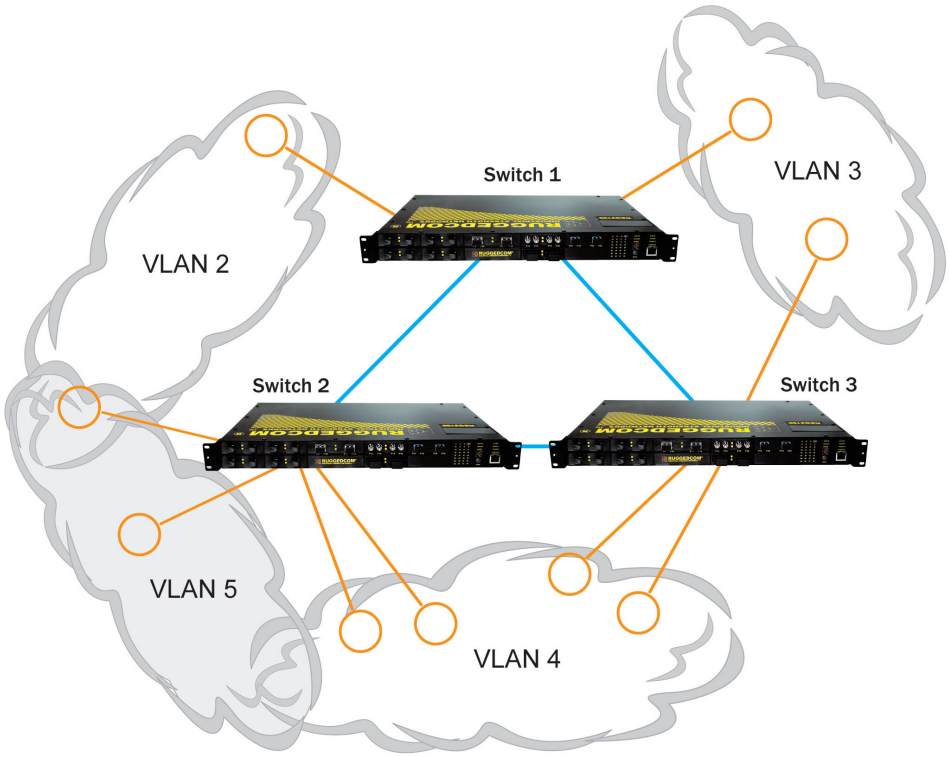

Figure 29.2. Multiple Overlapping VLANs

29.2.2. Administrative Convenience

VLANs enable equipment moves to be handled by software reconfiguration instead of by physical cable management. When a host's physical location is changed, its connection point is often changed as well. With VLANs, the host's VLAN membership and priority are simply copied to the new port.

29.2.3. Reduced Hardware

Without VLANs, traffic domain isolation requires using separate bridges for separate networks. VLANs eliminate the need for separate bridges.

The number of network hosts may often be reduced. Often, a server is assigned to provide services for independent networks. These hosts may be replaced by a single, multi-homed host supporting each network on its own VLAN. This host can perform routing between VLANs.

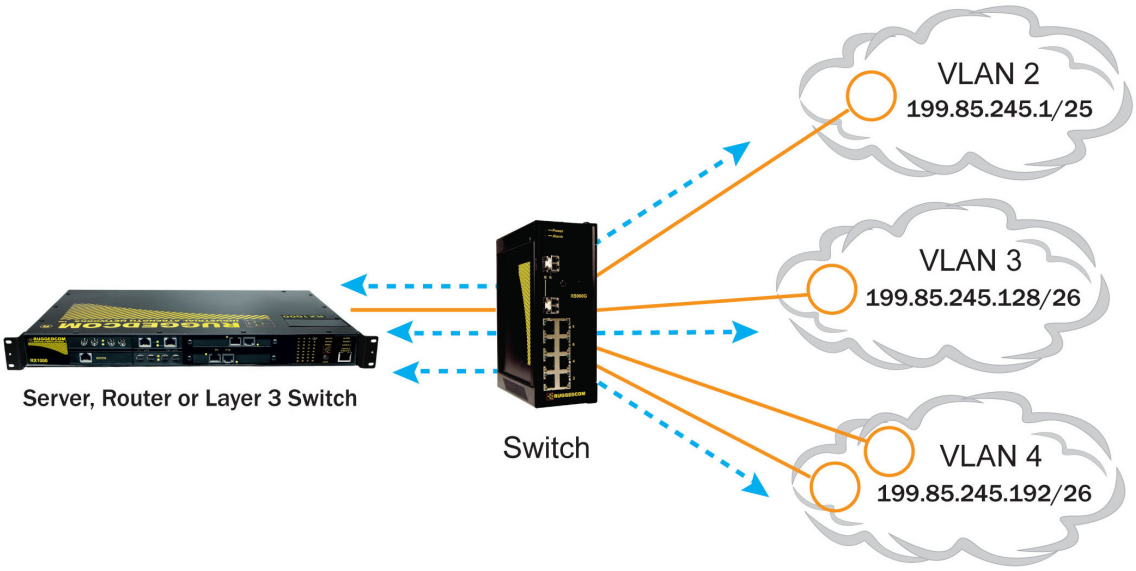

Figure 29.3. Inter-VLAN Communications

29.3. VLAN Configuration

The Virtual LANs menu and the Internal VLAN Range form are accessible from the main menu under *switch/vlans*.

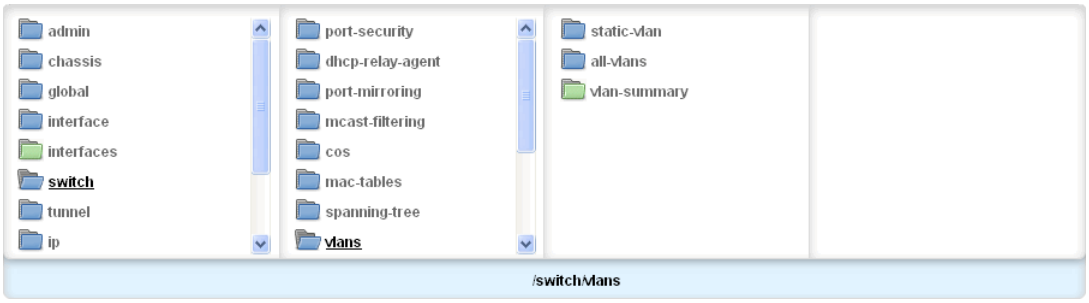

Figure 29.4. Virtual LANs menu

Figure 29.5. Internal VLAN Range form

29.3.1. Static VLANs

If static VLANs have been configured, the Static VLAN table will be displayed under *switch/vlans/static-vlan*. To display the forms, navigate to *switch/vlans/static-vlan/{number}*.

Figure 29.6. Static VLAN table

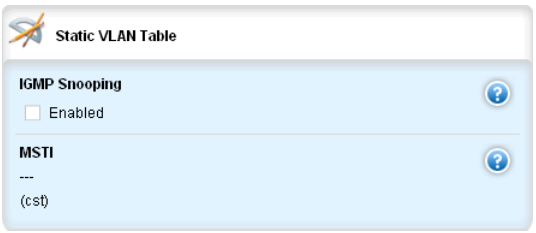

Figure 29.7. Static VLAN form

VLAN ID

Synopsis: *integer*

The VLAN Identifier is used to identify the VLAN in tagged Ethernet frames according to IEEE 802.1Q.

IGMP Snooping

Enables or disables IGMP Snooping on the VLAN.

MSTI

Synopsis: *string - the keyword { cst }*

Synopsis: *integer*

Default: *cst*

Only valid for Multiple Spanning Tree Protocol (MSTP) and has no effect, if MSTP is not used. The parameter specifies the Multiple Spanning Tree Instance (MSTI) the VLAN should be mapped to.

> *If IGMP Snooping is not enabled for the VLAN, both IGMP messages and multicast streams will be forwarded directly to all members of the VLAN. If any one member of the VLAN joins a multicast group then all members of the VLAN will receive the multicast traffic.*

29.3.2. Port VLAN Parameters

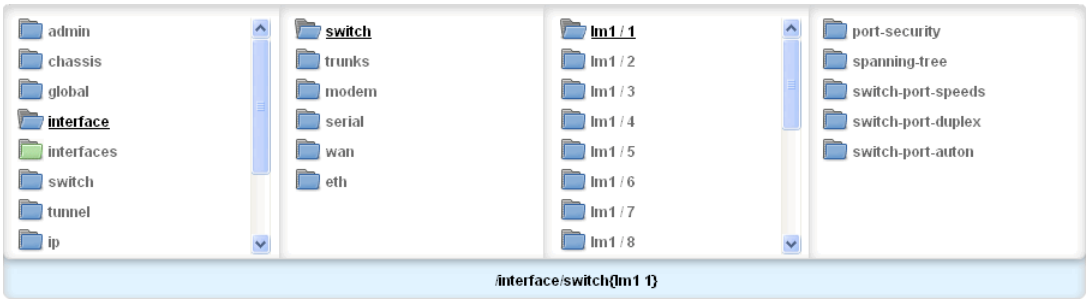

Figure 29.8. Switched Ethernet Ports submenu

The VLAN parameter forms can be accessed in two locations: *interface/switch/{line module}* (for example, lm1/1) or *interface/trunks/{number}*.

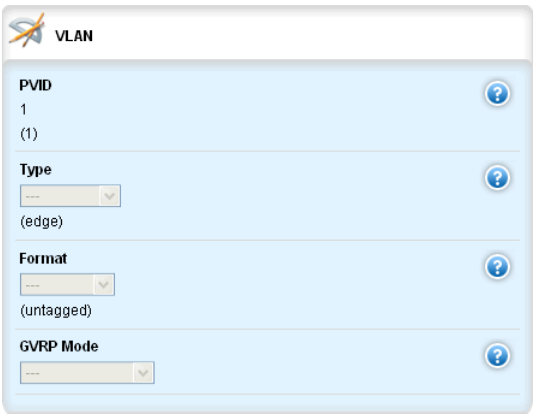

Figure 29.9. VLAN Parameters form

PVID

Synopsis: *integer*

Default: *1*

The Port VLAN Identifier specifies the VLAN ID associated with untagged (and 802.1p priority tagged) frames received on this port. Frames tagged with a non-zero VLAN ID will always be associated with the VLAN ID retrieved from the frame tag.

Type

Synopsis: *string - one of the following keywords { pvlanedge, trunk, edge }*

Default: *edge*

How the port determines its membership in VLANs. There are a few types of ports:

- EDGE : the port is only a member of one VLAN (its native VLAN specified by the 'PVID' parameter).
- PVLANEdge : the port does not forward traffic to other PVLANedge ports within the same VLAN.
- TRUNK : the port is automatically a member of all configured VLANs. Frames transmitted out of the port on all VLANs except the port's native VLAN will be always tagged. It can also be configured to use GVRP for automatic VLAN configuration.

Format

Synopsis: *string - one of the following keywords { tagged, untagged }*

Default: *untagged*

Whether frames transmitted out of the port on its native VLAN (specified by the 'PVID' parameter) will be tagged or untagged.

GVRP Mode

Synopsis: *string - one of the following keywords { learn_advertise, advertise_only }*

GVRP (Generic VLAN Registration Protocol) operation on the port. There are several GVRP operation modes:

- DISABLED : the port is not capable of any GVRP processing.
- ADVERTISE ONLY : the port will declare all VLANs existing in the switch (configured or learned) but will not learn any VLANs.
- ADVERTISE and LEARN : the port will declare all VLANs existing in the switch (configured or learned) and can dynamically learn VLANs.

29.3.3. VLAN Summary

There are actually three ways that a VLAN can be created in the switch:

Explicit

A VLAN is explicitly configured in the Static VLANs list.

Implicit

A VLAN ID is a parameter required for different feature configurations (e.g. Port VLAN Parameters, Static MAC Addresses, IP Interface Type and ID). When such a parameter is set to some VLAN ID value, an appropriate VLAN is automatically created, if it does not yet exist.

Dynamic

A VLAN learned through GVRP.

Not explicitly created VLAN is always created with IGMP Snooping disabled. If it is desirable for IGMP to be used on that VLAN, it should be created as a Static VLAN with IGMP enabled.

All VLANs, regardless of the way they were created, are shown in the VLAN Summary.

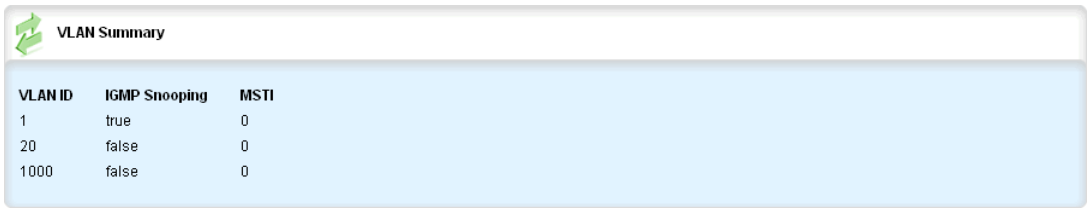

Figure 29.10. VLAN Summary table

To display the VLAN Summary table, navigate to *switch/vlans/vlan-summary*.

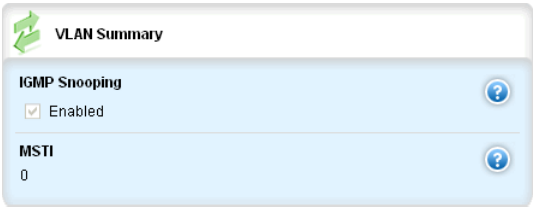

Figure 29.11. VLAN Summary form

VLAN ID

Synopsis: *integer*

The VLAN Identifier is used to identify the VLAN in tagged Ethernet frames according to IEEE 802.1Q.

IGMP Snooping

Synopsis: *boolean*

Enables/disables IGMP-Snooping.

MSTI

Synopsis: *integer*

The assigned MSTP Instance ID.

To display the VLAN Summary form, navigate to *switch/vlans/vlan-summary/{number}*.

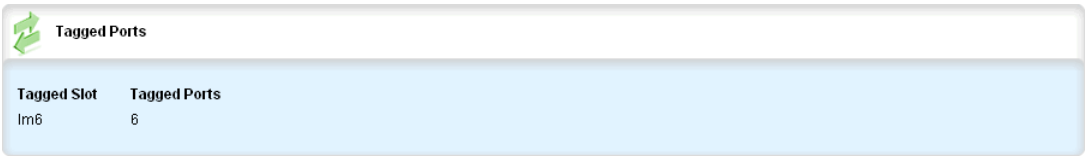

Figure 29.12. Tagged Ports table

Tagged-ports and untagged-ports can be viewed. To display the Tagged Ports table, navigate to *switch/vlans/vlan-summary/{number}/tagged-ports*.

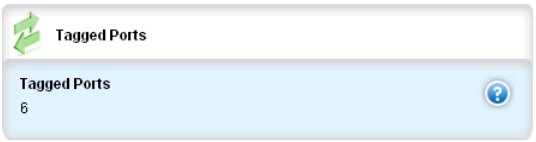

Figure 29.13. Tagged Ports form

To display the Tagged Ports form, navigate to *switch/vlans/vlan-summary/{number}/tagged-ports/{line module}*.

Tagged Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Tagged Ports

Synopsis: *A string*

The selected ports on the module installed in the indicated slot.

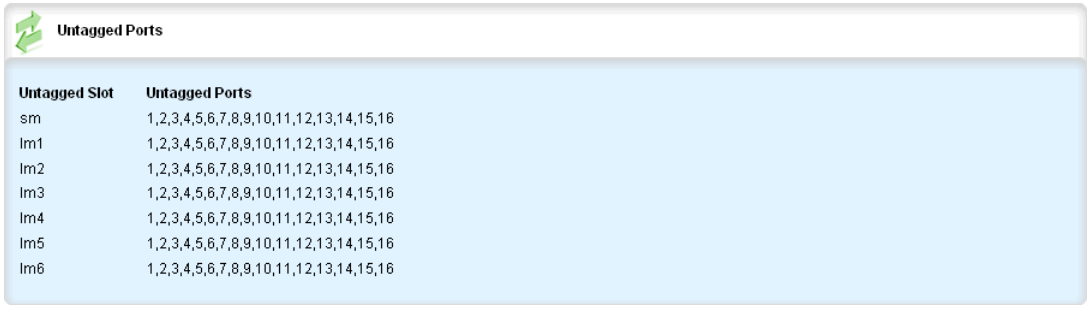

Figure 29.14. Untagged Ports table

To display the Tagged Ports table, navigate to *switch/vlans/vlan-summary/{number}/untagged-ports*.

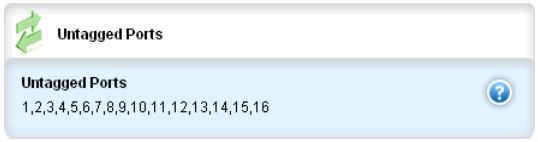

Figure 29.15. Untagged Ports form

To display the Tagged Ports form, navigate to *switch/vlans/vlan-summary/{number}/untagged-ports/{line module}*.

Untagged Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Untagged Ports

Synopsis: *A string*

The selected ports on the module installed in the indicated slot.

Figure 29.16. All VLANs table

To display the All VLANs table, navigate to *switch/vlans/all-vlans*.

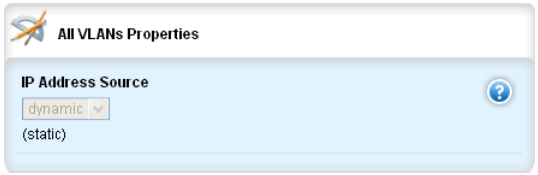

Figure 29.17. All VLANs Properties form

To display the All VLANs Properties form, navigate to *switch/vlans/all-vlans/{number)*.

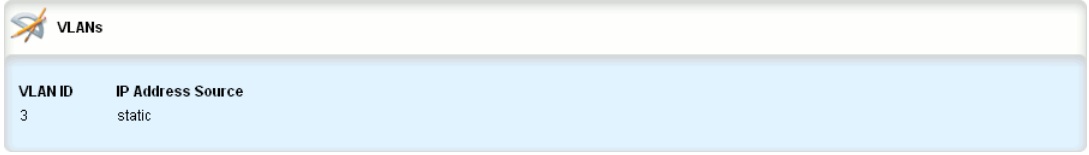

Figure 29.18. VLANs table

To display the VLANs table, navigate to *interface/eth/{line module}/vlan*.

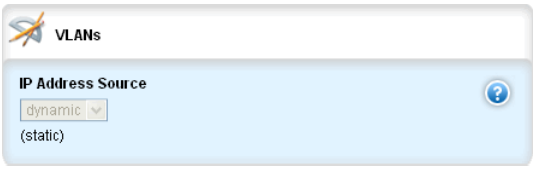

Figure 29.19. VLANs form

To display the VLANs form, navigate to *interface/eth/{line module}/vlan/{number}*.

VLAN ID

Synopsis: *integer*

VLAN ID for this routable logical interface

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP. The DYNAMIC option is a common case of a dynamically assigned IP address. It switches between BOOTP and DHCP until it gets the response from the relevant server. It must be static for non-management interfaces

on-demand

This interface is up or down on demand of link fail over.

29.3.4. Forbidden Ports

Figure 29.20. Forbidden Ports

If forbidden ports are configured, display the Forbidden Ports form by navigating to *switch/vlans/static-vlan/{number}/forbidden-ports*.

Slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The selected ports on the module installed in the indicated slot.

29.4. Troubleshooting

- *I don't need VLANs at all. How do I turn them off?*

  Simply leave all ports set to type "Edge" and leave the native VLAN set to 1. This is the default configuration for the switch.

- *I have added two VLANs: 2 and 3. I made a number of ports members of these VLANs. Now I need some of the devices in one VLAN to send messages to some devices in the other VLAN.*

  If the devices need to communicate at the physical address layer, they must be members of the same VLAN. If they can communicate in a Layer 3 fashion (i.e. using a protocol such as IP or IPX), you

can use a router. The router will treat each VLAN as a separate interface, which will have its own associated IP address space.

30. Network Discovery

Figure 30.1. Net-discovery menu

The Net-discovery menu is accessible from the main menu under *switch*. The path to this menu is *switch/net-discovery*.

ROX™ supports LLDP (the Link Layer Discovery Protocol), a Layer 2 protocol for automated network discovery.

LLDP is an IEEE standard protocol, IEEE 802.11AB, which allows a networked device to advertise its own basic networking capabilities and configuration. ROX™ is capable of advertising and collecting network information via LLDP. LLDP functionality in ROX™ includes the ability to:

- Enable or disable LLDP reception and transmission per port or for the whole device.
- View LLDP statistics.
- View 'neighbor' information.
- Report LLDP neighbor information via SNMP.

30.1. LLDP Operation

The IEEE standard, 802.1AB Link Layer Discovery Protocol (LLDP), describes a protocol that can simplify the troubleshooting of complex networks and can be used by Network Management Systems (NMS) to obtain and monitor detailed information about a network's topology. LLDP data are made available via SNMP (through support of LLDP-MIB).

LLDP allows a networked device to discover its neighbors across connected network links using a standard mechanism. Devices that support LLDP are able to advertise information about themselves, including their capabilities, configuration, interconnections, and identifying information.

LLDP agent operation is typically implemented as two modules: the LLDP transmit module and LLDP receive module. The LLDP transmit module, when enabled, sends the local device's information at regular intervals, in 802.1AB standard format. Whenever the transmit module is disabled, it transmits an LLDPDU (LLDP data unit) with a time-to-live (TTL) TLV containing "0" in the information field. This enables remote devices to remove the information associated with the local device in their databases. The LLDP receive module, when enabled, receives remote devices' information and updates its LLDP database of remote systems. When new or updated information is received, the receive module initiates a timer for the valid duration indicated by the TTL TLV in the received LLDPDU. A remote system's information is removed from the database when an LLDPDU is received from it with TTL TLV containing "0" in its information field.

LLDP is implemented to keep a record of only one device per Ethernet port. Therefore, if there are multiple devices sending LLDP information to a switch port on which LLDP is enabled, information about the neighbor on that port will change constantly.

30.2. LLDP Parameters

Figure 30.2. Net-discovery LLDP menu

The Net-discovery LLDP menu is accessible from the main menu under *switch*. The path to this menu is *switch/net-discovery/lldp*. The LLDP form, LLDP Global Statistics Form and LLDP Local System form appear on the same screen as the menu.

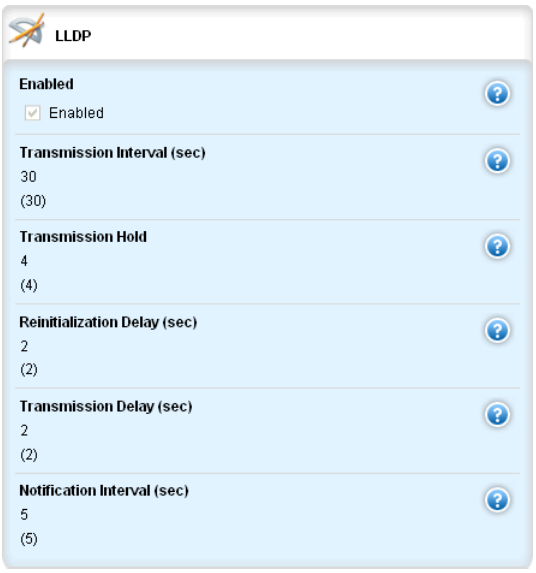

Figure 30.3. LLDP form

This form is used to configure the Network Discovery Protocol (LLDP).

Enabled

Synopsis: *boolean*

Default: *true*

Enables the LLDP protocol. Note that LLDP is enabled on a port when LLDP is enabled globally and along with enabling per port setting in Port LLDP Parameters menu.

Transmission Interval (sec)

Synopsis: *integer*

Default: *30*

The interval at which LLDP frames are transmitted on behalf of this LLDP agent.

Transmission Hold

Synopsis: *integer*

Default: *4*

The multiplier of the Tx Interval parameter that determines the actual time-to-live (TTL) value used in an LLDPDU. The actual TTL value can be expressed by the following formula: TTL = MIN(65535, (Tx Interval * Tx Hold))

Reinitialization Delay (sec)

Synopsis: *integer*

Default: *2*

The delay in seconds from when the value of the Admin Status parameter of a particular port becomes 'Disbled' until re-initialization will be attempted.

Transmission Delay (sec)

Synopsis: *integer*

Default: *2*

The delay in seconds between successive LLDP frame transmissions initiated by the value or status changed. The recommended value is set by the following formula: 1 is less than or equal to txDelay less than or equal to (0.25 * Tx Interval)

Notification Interval (sec)

Synopsis: *integer*

Default: *5*

Controls transmission of LLDP traps. The agent must not generate more than one trap in an indicated period.

Figure 30.4. LLDP Global Statistics form

Inserts

Synopsis: *unsigned integer*

The number of times an entry was inserted into the LLDP Neighbor Information Table.

Deletes

Synopsis: *unsigned integer*

The number of times an entry was deleted from the LLDP Neighbor Information Table.

Drops

Synopsis: *unsigned integer*

The number of times an entry was deleted from the LLDP Neighbor Information Table because the information timeliness interval has expired.

Ageouts

Synopsis: *unsigned integer*

The number of all TLVs discarded.

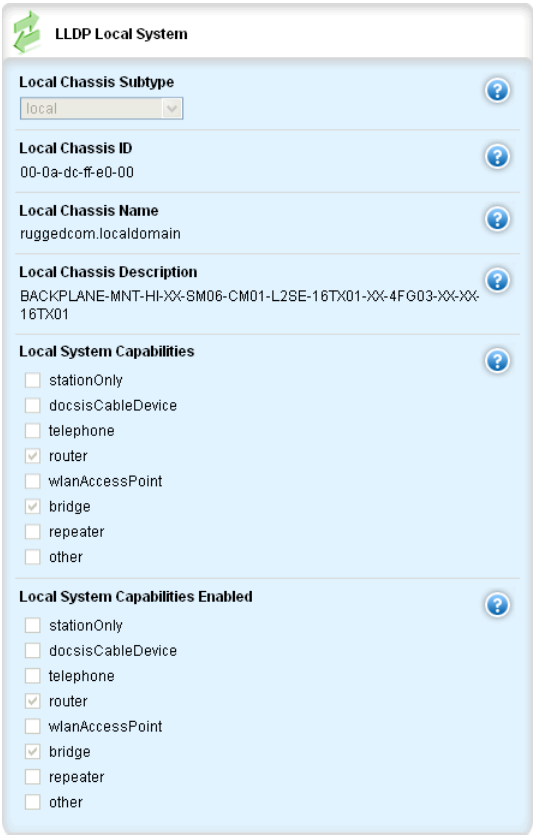

Figure 30.5. LLDP Local System form

The LLDP local system form provides access to the local host's information that is being set to remote LLDP-enabled devices.

Local Chassis Subtype

Synopsis: *string - one of the following keywords { local, interfaceName, networkAddress, macAddress, portComponent, interfaceAlias, chassisComponent }*

local-chassis-subtype

Local Chassis ID

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

local-chassis-id

Local Chassis Name

Synopsis: *A string*

local-system-name

Local Chassis Description

Synopsis: *A string*

local-system-desc

Local System Capabilities

local-system-caps

Local System Capabilities Enabled

local-system-caps-enabled

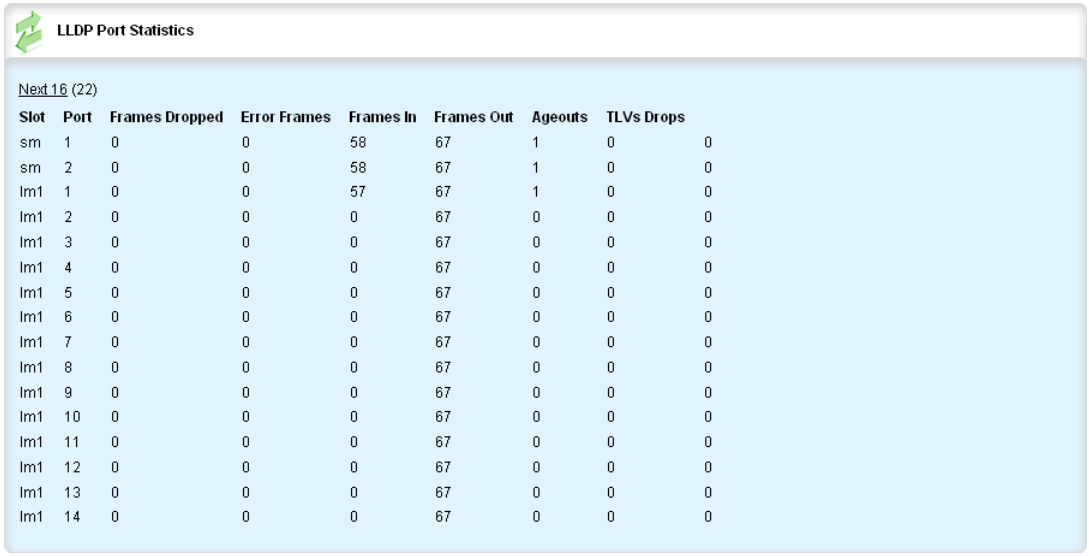

LLDP Port Statistics

Next 16 (22)

Slot	Port	Frames Dropped	Error Frames	Frames In	Frames Out	Ageouts	TLVs Drops	
sm	1	0	0	58	67	1	0	0
sm	2	0	0	58	67	1	0	0
lm1	1	0	0	57	67	1	0	0
lm1	2	0	0	0	67	0	0	0
lm1	3	0	0	0	67	0	0	0
lm1	4	0	0	0	67	0	0	0
lm1	5	0	0	0	67	0	0	0
lm1	6	0	0	0	67	0	0	0
lm1	7	0	0	0	67	0	0	0
lm1	8	0	0	0	67	0	0	0
lm1	9	0	0	0	67	0	0	0
lm1	10	0	0	0	67	0	0	0
lm1	11	0	0	0	67	0	0	0
lm1	12	0	0	0	67	0	0	0
lm1	13	0	0	0	67	0	0	0
lm1	14	0	0	0	67	0	0	0

Figure 30.6. LLDP Port Statistics table

The LLDP Port Statistics table allows you to view port LLDP statistics

The path to the LLDP Port Statistics table is *switch/net-discovery/lldp/port-lldp-stats*.

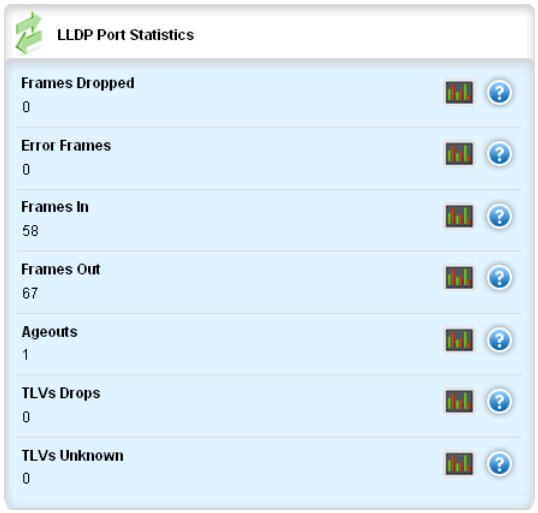

Figure 30.7. LLDP Port Statistics form

The path to the LLDP Port Statistics form is *switch/net-discovery/lldp/port-lldp-stats* and then clicking on one of the linked submenus (for example, sm/1).

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the module.

Frames Dropped

Synopsis: *unsigned integer*

A counter of all LLDP frames discarded

Error Frames

Synopsis: *unsigned integer*

A counter of all LLDPDUs received with detectable errors

Frames In

Synopsis: *unsigned integer*

A counter of all LLDPDUs received

Frames Out

Synopsis: *unsigned integer*

A counter of all LLDPDUs transmitted

Ageouts

Synopsis: *unsigned integer*

A counter of the times that a neighbor's information has been deleted from the LLDP remote system MIB because the txinfoTTL timer has expired

TLVs Drops

Synopsis: *unsigned integer*

A counter of all TLVs discarded

TLVs Unknown

Synopsis: *unsigned integer*

A counter of all TLVs received on the port that are not recognized by the LLDP local agent

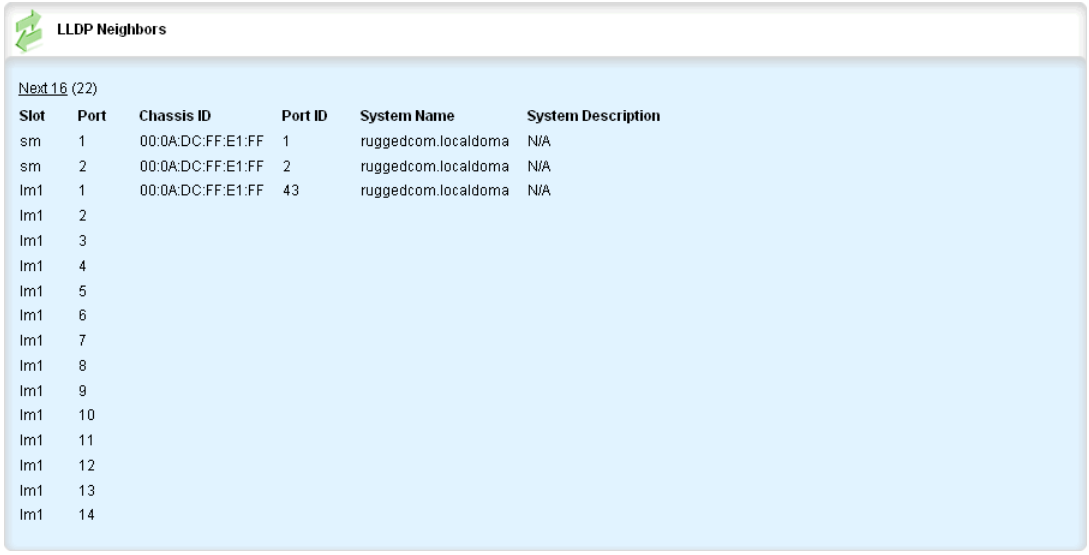

Figure 30.8. LLDP Neighbors table

The LLDP Neighbors table allows you to view LLDP neighbors by port.

The path to the LLDP Neighbors table is *switch/net-discovery/lldp/port-lldp-neighbors*.

Figure 30.9. LLDP Neighbors form

The path to the LLDP Neighbors form is *switch/net-discovery/lldp/port-lldp-neighbors* and then clicking on one of the linked submenus (for example, sm/1).

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the module.

Chassis ID

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

Chassis ID information received from a remote LLDP agent.

Port ID

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

Port ID (MAC) information received from a remote LLDP agent.

System Name

Synopsis: *A string*

System name information received from a remote LLDP agent

System Description

Synopsis: *A string*

System descriptor information received from a remote LLDP agent.

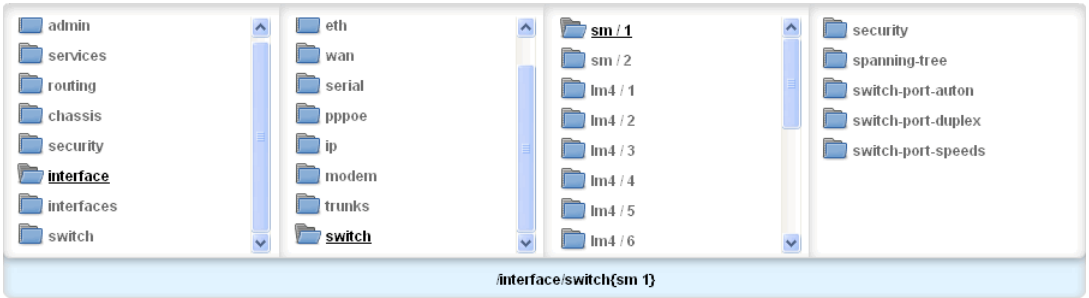

Figure 30.10. LLDP submenu

The LLDP submenu is accessible from the main menu under *interface*. The path to this menu is *interface/switch* and then clicking on one of the linked submenus (for example, sm/1).

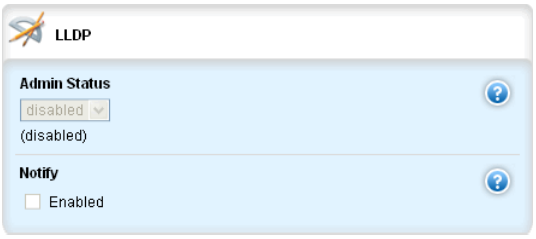

Figure 30.11. LLDP form

Admin Status

Synopsis: *string - one of the following keywords { no-lldp, rx-tx, rx-only, tx-only }*

Default: *rx-tx*

- no-lldp : The local LLDP agent can neither transmit nor receive LLDP frames.
- rxTx : The local LLDP agent can both transmit and receive LLDP frames through the port.
- txOnly : The local LLDP agent can only transmit LLDP frames.
- rxOnly : The local LLDP agent can only receive LLDP frames.

Notify

Disabling notifications will prevent sending notifications and generating alarms for a particular interface from the LLDP agent. If a domain name is not specified here, the router will attempt to extract this information from the host addresses.

Part III. Routing and Security

This part describes routing and network security:

Routing Overview	Chapter 31, *ROX™ Routing Overview*
Layer 3 Switching	Chapter 32, *Layer 3 Switching*
Tunnelling	Chapter 33, *Tunnelling*
Dynamic Routing	Chapter 34, *Dynamic Routing*
Static Routing	Chapter 35, *Static Routing*
Routing Status	Chapter 36, *Routing Status*
Multicast Routing	Chapter 37, *Multicast Routing*
Firewall	Chapter 38, *Firewall*
Traffic Control	Chapter 39, *Traffic Control*
VRRP	Chapter 40, *VRRP*
LinkFailover	Chapter 41, *Link Failover*

31. ROX™ Routing Overview

This section provides an overview of IP routing in ROX™. This section describes how ROX™ configures physical Ethernet ports, and how ROX™ switches and routes IP packets.

31.1. IP Routing in ROX™

ROX™ supports multiple interface types and can perform Ethernet Switching (Layer 2), IP switching (Layer 3), and Software routing depending on the type of physical interface. Almost all interface types support IP addresses in ROX™.

31.2. Physical Ethernet Port Types in ROX™

ROX™ supports two types of Ethernet ports: *switch ports* and *router (Ethe) ports*:

- A *switch port* is one in which packet switching is performed in the underlying hardware. By default, all switch ports belong to VLAN 1.
- A *router port* is one in which packet switching is managed entirely in software. Router ports are usually independent of other ports or interfaces and are managed as such. Hardware acceleration cannot be performed on router ports.

On the RX1500-series and RX5000 platforms, all Ethernet ports except for cm/1 are switch ports. On the RX1000-series platforms, all Ethernet ports are router ports.

31.3. Routing

31.3.1. Using VLAN Interfaces for Routing and Layer 3 Switching

When a VLAN interface is configured with an IP address, ROX™ hosts the IP on the LAN segment. This means that any devices belonging to that VLAN can reach the IP address, provided that the devices' IP addresses belong to the same subnet.

For example, assume that lm1/1, lm1/2, and lm1/5 belong to VLAN 100 and the attached devices have the IP addresses 192.168.100.1/24, 192.168.100.2/24, and 192.168.100.3/24. At this point, the ROX™ device does not have an IP address for VLAN 100, but it has an IP address for VLAN 1. In this example, devices attached to lm1/1, lm1/2, and lm1/5 can communicate among themselves, but their LAN segment is essentially isolated and they cannot reach anywhere else.

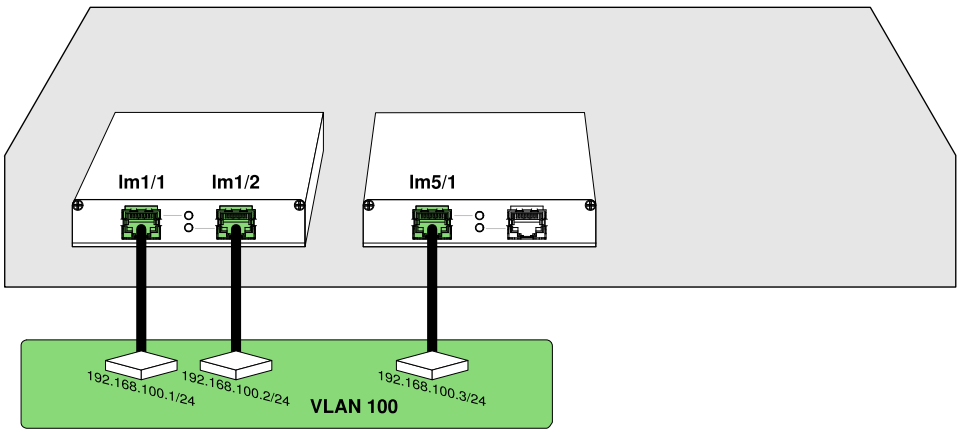

Figure 31.1. Three interfaces on an isolated VLAN

Whenever you create a VLAN in ROX™, a corresponding interface with no IP address is also created. This interface is accessible to devices attached to that VLAN.

Continuing with the example above, an IP interface with the name Switch.0100 is created when you create VLAN 100. Providing an IP address to this interface makes the ROX™ system accessible to the devices on VLAN 100. For example, assigning Switch.0100 the IP address 192.168.100.10/24 makes the ROX™ device accessible to the devices attached to lm1/1, lm1/2, and lm1/5. Note that all IP addresses belong to the same subnet.

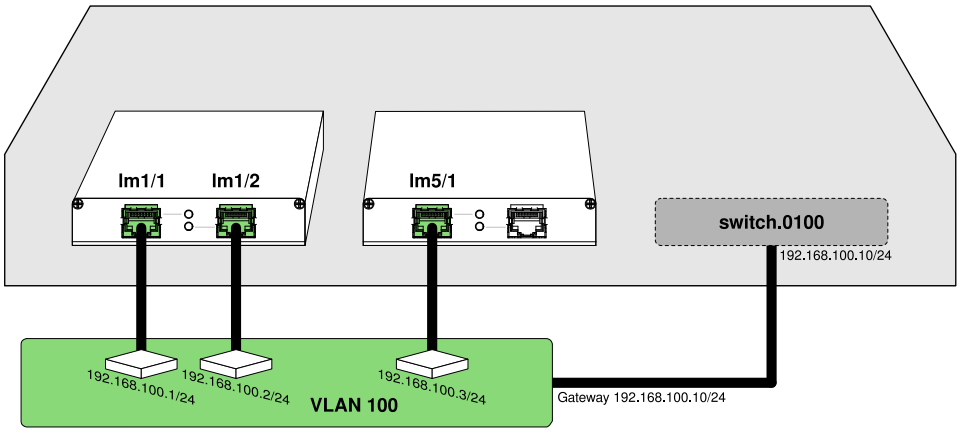

Figure 31.2. VLAN connected to ROX device through switch.0100

If the ROX™ device is to be used as a router to route between all attached networks, switch ports can be made to behave liks a routed port. For more information, see Section 31.3.2, “Routing IP Packets”.

31.3.2. Routing IP Packets

When the devices belonging to VLAN 100 need to reach another network, they can use Switch.0100 as the gateway. At that point, the ROX™ device routes the traffic. If the traffic volume to be routed is high enough, then Layer 3 switching will start (provided that this feature is available). Note that the devices attached to lm1/1, lm1/2, and lm1/5 must use Switch.0100’s IP address as the gateway.

Routed Ports

To make a particular port behave like a router port, do the following:

- disable the spanning tree protocol on the port.
- assign the port to a unique VLAN.
- assign the VLAN interface an IP address.

The VLAN is still an isolated LAN segment with an IP address, but it has only one physical port attached; thus, the port will behave as if it is a router port. For example, if the port lm2/3 belongs to VLAN 1760 and no other port belongs to that VLAN on the local system, then Switch.1760 can be directly reached only through lm2/3, and thus behaves as a router port.

Note that the port still belongs to a VLAN. If the port has to behave like a router interface, it should be configured as a VLAN edge port. Also note that trunk ports, by definition, belong to all VLANs, unless they are specifically configured otherwise.

32. Layer 3 Switching

32.1. Layer 3 Switching Fundamentals

32.1.1. What is a Layer 3 Switch?

A *switch* is an internetwork device that makes frame forwarding decisions in hardware. A *Layer 3 switch*, sometimes called a multilayer switch, is one which makes hardware-based decisions for IP packets as well as Layer 2 frames. Traditionally, routers are used to make routing decisions using software. A Layer 3 switch will make the same decisions in hardware, which means that packet forwarding will be much faster than in a conventional router.

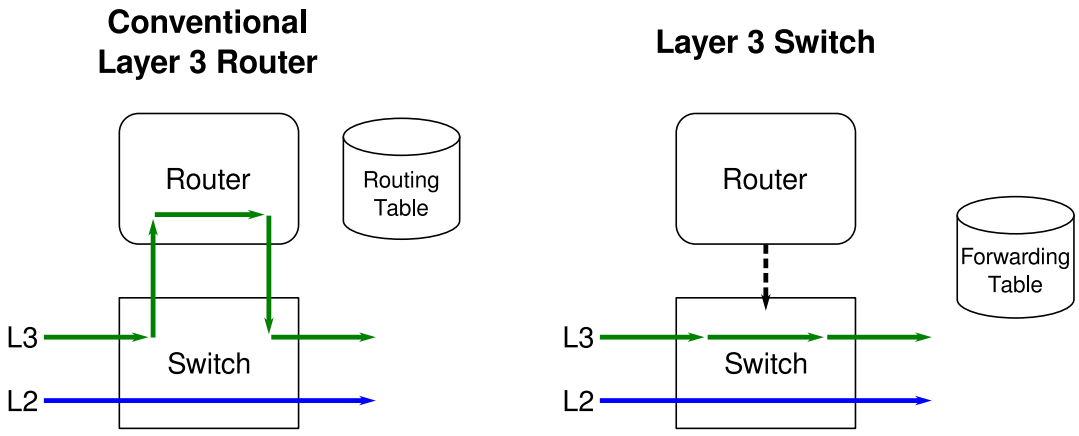

Figure 32.1. Layer 3 Switch

The RuggedBackbone™ Layer 3 Switch combines the rich feature set of a software router and the wire-speed performance of a Layer 3 switch. It offers flexible configuration, allowing you to control which routed traffic flows are hardware-accelerated and which flows are subject to software processing. This allows the device to satisfy sophisticated firewall rules, which are not normally not supported by Layer 3 switches.

32.1.2. Layer 3 Switch Forwarding table

To route a packet with a specific destination IP address, a router needs the following information:

- **Egress interface (subnet):** this information is stored in the router's Routing Table.

> *In a Layer 2 switched network segment, a VLAN constitutes an IP subnet.*

- **Next-hop gateway MAC address:** this information is stored in the router's ARP Table.

> *If the next hop is the destination subnet itself, then the destination host MAC address is required.*

A Layer 3 Switch uses the routing information listed above and translates it into Layer 3 switching rules. These rules are known as the *Layer 3 Switch Forwarding Information Base (FIB)* or the *Layer 3 Switch Forwarding Table*. A Layer 3 switching rule is actually a set of parameters identifying a traffic flow to be switched and determining how to perform the switching.

Layer 3 switching application-specific integrated circuits (ASICs) store Layer 3 switching rules in a Ternary Content Addressable Memory (TCAM) table. Layer 3 switching rules can be statically configured or dynamically learned (also known as *auto-learned*).

32.1.3. Static Layer 3 Switching Rules

When creating a static route through switch management, you can explicitly configure it to be hardware-accelerated. If hardware acceleration is selected, an appropriate Layer 3 switching rule is installed in the ASIC's TCAM and never ages out.

32.1.4. Dynamic Learning of Layer 3 Switching Rules

For static routes without hardware acceleration or for dynamic routes, Layer 3 switching rules can be dynamically learned based on software router and firewall decisions. For example, the Layer 3 switch can automatically decide to offload some flows from the router into the Layer 3 Forwarding Table.

After a certain amount of traffic for the same flow is successfully routed, the Layer 3 switching ASIC begins switching the rest of the packets belonging to the same flow. A flow is unidirectional traffic between two hosts. For example, the traffic from 192.168.10.1/24 TCP port 1789 to 192.168.20.1/24 TCP port 1623 is a flow. Traffic in the opposite direction constitutes another flow.

The RuggedBackbone™ Layer 3 Switch supports different modes of dynamic rule learning.

Flow-oriented learning is when the switch uses the following information to identify a traffic flow:

- Source IP address
- Destination IP address
- Protocol
- Source TCP/UDP port
- Destination TCP/UDP port

This learning method is more granular and requires more ASIC resources, but it provides more flexibility in firewall configuration as the rule takes the protocol and TCP/UDP port into consideration to make forwarding decisions.

Host-oriented learning is when the switch uses the following information to identify a traffic flow:

- Source IP address
- Destination IP address

This learning method provides less flexibility in firewall configuration, as the user can allow or disallow traffic between two hosts.

For unicast traffic, each flow constitutes one rule. For multicast routing, one multicast route may constitute several rules. For more information, see Section 32.1.6, "Layer 3 Multicast Switching".

The Layer 3 switch continuously monitors activity (this is, the presence of traffic) for dynamically learned rules. Because of this, dynamically learned rules may be removed after a configurable time due to inactivity.

32.1.5. Layer 3 Switch ARP table

A router needs to know the destination host or next-hop gateway MAC address for it to forward a packet on the other subnet. Therefore, software maintains an ARP (Address Resolution Protocol) table that maps IP addresses to MAC addresses. The same information is also needed by the Layer 3 switching ASIC when it switches IP packets between subnets.

The destination or gateway MAC address is usually obtained through ARP. However, ARP entries can also be statically configured in the Layer 3 Switch so that they do not time out. When configuring a static ARP entry, if no value is entered for the MAC Address parameter, the address is automatically resolved through ARP and then saved statically. This is preserved across reboots of the device.

For a static Layer 3 switching rule, the destination MAC address for the rule is always resolved, and is also saved statically.

32.1.6. Layer 3 Multicast Switching

Some RuggedCom Layer 3 Switch models do not have full multicast Layer 3 switching capability and only support multicast cross-VLAN Layer 2 switching. Multicast cross-VLAN Layer 2 switching differs from the normal multicast Layer 3 switching in the following ways:

- Packet modification is not done. That is, the source MAC address and TTL values in forwarded packets do not change. This should not be a problem in most cases, but it should be taken into consideration.
- Cross-VLAN Layer 2 switching is less efficient in ASIC resource utilization and packet latency.
- Separate TCAM table entries are required for each egress VLAN in the multicast switching rule. For example, a multicast stream ingressing VLAN 1 and egressing VLAN 2 and VLAN 3 requires two TCAM table entries: one for VLAN 2 and one for VLAN 3.
- Supported bandwidth depends on the rule. Multicast traffic potentially has multiple egress VLANs, and the total utilized ASIC bandwidth is the ingress bandwidth multiplied by the number of ingress and egress VLANs. For example, a 256Mbps multicast stream ingressing VLAN 1 and egressing VLANs 2 and 3 requires 768Mbps (256Mbps × 3) of ASIC bandwidth.
- If a multicast packet should be forwarded to multiple egress VLANs, it egresses those VLANs sequentially rather than concurrently. This means that the packet will experience different latency for each egress VLAN.

32.1.7. Size of the Layer 3 Switch Forwarding Table

The routing table in a software router is limited only by the amount of available memory; its size can be virtually unlimited. However, the size of the TCAM in Layer 3 switching ASICs is significantly limited and may not be sufficient to accommodate all Layer 3 switching rules. If the TCAM is full and a new static rule is created, the new rule replaces some dynamically learned rule. If all of the rules in the TCAM are static, then the new static rule is rejected.

32.1.8. Interaction with the Firewall

If security is a concern and you use a firewall in a Layer 3 Switch, it is important to understand how the Layer 3 switch interacts with the firewall.

A software router always works in agreement with a firewall so that firewall rules are always applied. However, in a Layer 3 Switch, if a switching rule is set in the switching ASIC (for example, due to a statically configured route), the ASIC switches all the traffic matching the rule before the firewall inspects the traffic.

Layer 3 switch ASICs are somewhat limited in how switching rules can be defined. These limitations do not allow configuring arbitrary firewall rules directly in the Layer 3 switch hardware. For sophisticated firewall rules, the firewall has to be implemented in software and the Layer 3 Switch must not switch traffic that is subject to firewall processing.

Whenever a change is made to the firewall configuration, some of the dynamically learned Layer 3 switching rules might “conflict” with the new firewall configuration. To resolve potential conflicts, dynamically learned Layer 3 switching rules are flushed upon any changes to the firewall configuration. The dynamically learned Layer 3 switching rules then have to be re-learned while the new firewall rules are applied.

For statically configured Layer 3 switching rules, take care to avoid conflicts between Layer 3 switching and the firewall. It should be understood that static Layer 3 switching rules always take precedence. Therefore, you must thoroughly examine the switch configuration for potential conflicts with the firewall.

32.1.9. Sample Use Case

Consider the network illustrated below. The switch connecting all of these networks is a RuggedBackbone™ Layer 3 switch.

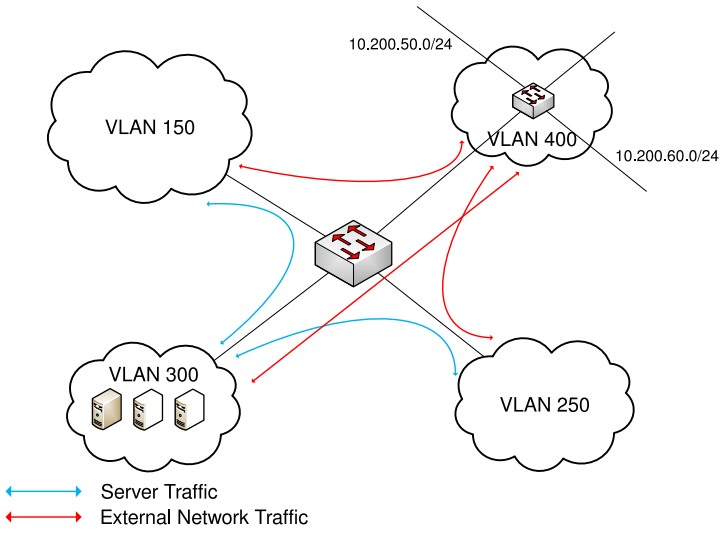

Figure 32.2. Layer 3 Switch Use Case

Assume the following:

- VLAN 150 and VLAN 250 have approximately 200 devices each.
- VLAN 300 is a server farm with RuggedNMS and some other servers polling these devices on a near-constant basis. The IP address for Server 1 is 172.30.30.10. The IP address for Server 2 is 172.30.30.20
- VLAN 400 connects the Layer 3 switch to an external network via a gateway with the IP address 172.30.40.2.
- Devices and servers in VLAN 150, 250 and 300 should only be able to reach 2 networks – 10.200.50.0/24 and 10.200.60.0/24 – for management purposes.
- The 400 devices in VLAN 150 and VLAN 250 receive IP multicast data from the external network (VLAN 400) at the address 227.100.20.100.
- Servers in VLAN 300 receive IP multicast data from the external network (VLAN 400) at the address 227.100.250.250.
- No firewall is used in this use case.

Assume all Layer 2 switching-related configuration has been done; that is, the VLANs have been created with proper port assignment. In this example, no specific Layer 3 switch configuration is explicitly required. As long as a software router configuration is sufficient for the above requirements, the Layer 3 Switch will be able to function through auto-learning switching rules. However, the Layer 3 switching configuration described in the following sections is recommended for better utilization of CPU and Layer 3 switching ASIC resources:

- Section 32.1.9.1, “Setting up Unicast Routes”
- Section 32.1.9.2, “Setting up Multicast Routing”
- Section 32.1.9.3, “Configuring Static Layer 3 Switching Rules for Servers”
- Section 32.1.9.4, “Configuring Static ARP Table Entries for Servers”
- Section 32.1.9.5, “Layer 3 Switching Rules for Devices in VLANs 150 and 250”

32.1.9.1. Setting up Unicast Routes

Because this use case only requires that the devices to be able to reach two networks, static routes can be used and can be hardware-accelerated.

- Create a static route in *routing/static/ipv4/route* and enter the network 10.200.50.0/24.
- Set *Hw-accelerate* to **Enabled**

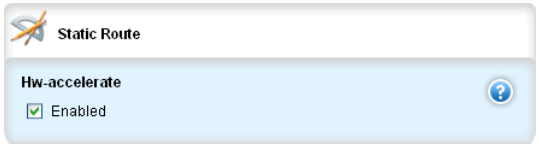

Figure 32.3. Hardware Acceleration Enabled

- Set the *via* parameter to the gateway's IP address (172.30.40.2).

This configuration creates the following:

- A Layer 3 Switch ARP Table entry specifying the gateway's resolved MAC address.
- A Layer 3 switching rule that switches any traffic destined for the 10.200.50.0/24 network.

The configuration can be verified under *switch/layer3-switching/arp-table* and *switch/layer3-switching/rules-summary*. Do the same for the 10.200.60.0/24 network.

> *Even if Hw-accelerate is not enabled, Layer 3 switching is still performed, but all switching rules for traffic destined to the just-configured subnets will have to be auto-learned on a per-flow basis, thus utilizing greater ASIC resources.*

32.1.9.2. Setting up Multicast Routing

Configure static multicast IP groups or VLANs 150 and 250.

- Create a static multicast route in routing/multicast/static/mcast-groups for address 227.100.20.100. Specify the source IP and ingress interface Switch.0400.
- Set the *Hw-accelerate* option to **Enabled**.

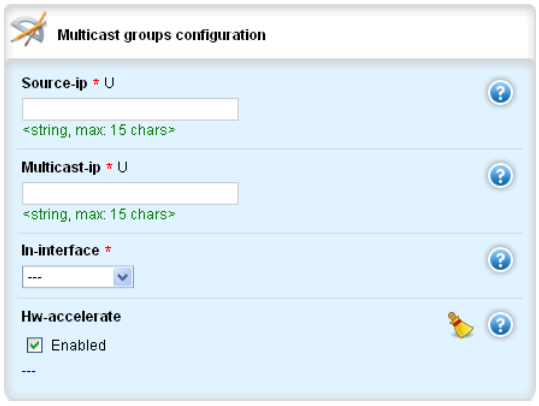

Figure 32.4. Hardware Acceleration Enabled

- Add Switch.0150 and Switch.0250 egress interfaces.

For servers:

- Create a static multicast route in *routing/multicast/static/mcast-groups* for address 227.100.250.250. Specify the source IP and ingress interface Switch.0400.
- Set *Hw-accelerate* to **Enabled**.

- Add egress interface Switch.0300.

This configuration creates Layer 3 switching rules which can be verified in *switch/layer3-switching/rules-summary*.

> *Even if Hw-accelerate is not enabled, Layer 3 switching is still performed, but all switching rules for the multicast streams will have to be auto-learned.*

32.1.9.3. Configuring Static Layer 3 Switching Rules for Servers

We can also create static Layer 3 switching rules for traffic destined to the servers so that those rules do not have to be auto-learned.

- Create a static route in *routing/static/ipv4/route* and enter the server's IP address (172.30.30.10/32).
- Set *Hw-accelerate* to **Enable**
- Set the *via* parameter to the same IP address (172.30.30.10).

Repeat these steps for each server.

32.1.9.4. Configuring Static ARP Table Entries for Servers

Even if configuring static rules for the servers is not desired, certain optimization is still possible. Each server in the server farm would be polling device IP addresses one after the other in order. Given that each server would always be talking to at least one device, we could create static ARP entries for the servers in the Layer 3 Switch ARP Table. This would keep the servers' MAC addresses always resolved so that the switch will not have to repeatedly resolve them when the addresses age out from the router's ARP table.

To create an ARP entry for a given server, navigate to *switch/layer3-switching/arp-table*. Enter the IP address of the server and then provide the MAC address of the server along with VID=300. If the MAC address of the server is not known, then leave it unspecified and the switch will automatically resolve the address and save it statically. Do the same for each server.

32.1.9.5. Layer 3 Switching Rules for Devices in VLANs 150 and 250

As the number of devices in VLANs 150 and 250 is quite large, it is impractical to configure a static route (or a static ARP table entry) for each of them, as we have done for the servers. Because of this, Layer 3 switching rules for traffic destined to those devices will have to be auto-learned.

32.2. Configuring Layer 3 Switching

To display the Layer 3 Switching menu, navigate to *switch/layer3-switching*.

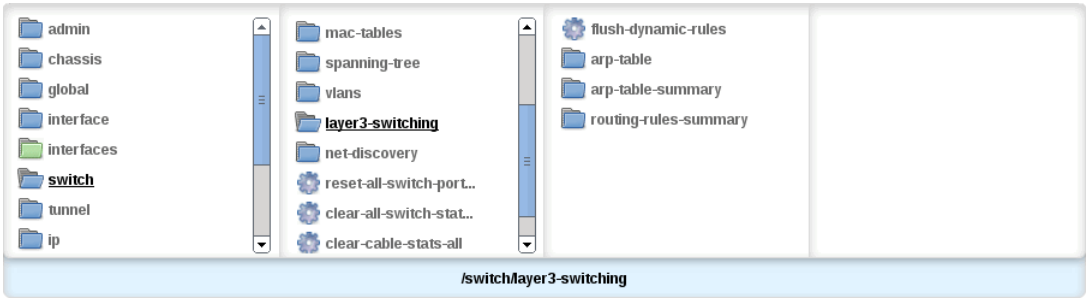

Figure 32.5. Layer 3 Switching menu

The Layer 3 Switching form on the menu page displays the configured Layer 3 switching settings.

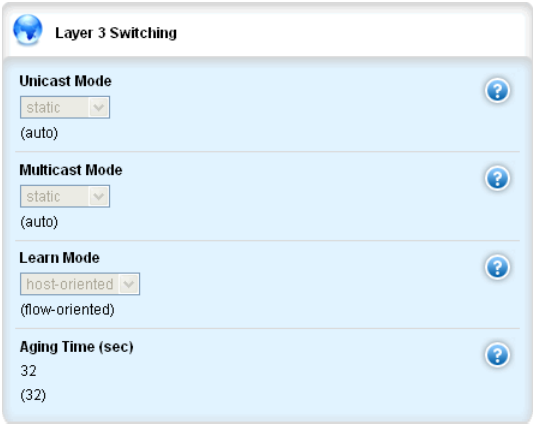

Figure 32.6. Layer 3 Switching form

To configure Layer 3 switching, do the following:

- set the Layer 3 switching settings. See Section 32.2.1, “Configuring Layer 3 Switching Settings”.
- create static ARP table entries. See Section 32.2.2, “Creating Static ARP Table Entries”.

After configuring Layer 3 switching, you can do the following:

- view static and dynamic ARP table entries. See Section 32.2.3, “Viewing Static and Dynamic ARP Table Entries”.
- view the routing rules summary. See Section 32.2.4, “Viewing Routing Rules”.
- flush dynamic hardware routing rules. See Section 32.2.5, “Flushing Dynamic Hardware Routing Rules”.

32.2.1. Configuring Layer 3 Switching Settings

To configure the Layer 3 switching settings:

- In edit mode, navigate to *switch/layer3-switching*.
- On the *Layer 3 Switching* form, set the Layer 3 switching parameters.
- Commit the changes.

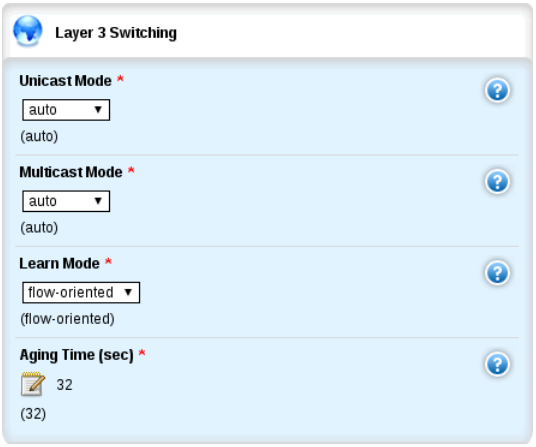

Figure 32.7. Layer 3 Switching form

Unicast Mode

Synopsis: *string - one of the following keywords { static, auto, disabled }*

Default: *auto*

- Disabled : Layer3 switching is disabled. The ability to disable routing hardware acceleration may be desired, for example, in a system with sophisticated firewall rules, which could not be supported by the Layer3 switching ASIC and would require software processing.
- Static : Only statically configured Layer3 switching rules will be used. This mode may be useful, for example, in a system with complex configuration where static routes do not conflict with a firewall, while traffic flows following dynamic routes have to be subject to sophisticated firewall filtering.
- Auto : Both statically configured and dynamically learned Layer3 switching rules will be used. In this mode, maximum routing hardware acceleration is utilized.

Multicast Mode

Synopsis: *string - one of the following keywords { static, auto, disabled }*

Default: *auto*

- Disabled : Layer3 switching is disabled. The ability to disable routing hardware acceleration may be desired, for example, in a system with sophisticated firewall rules, which could not be supported by the Layer3 switching ASIC and would require software processing.
- Static : Only statically configured Layer3 switching rules will be used. This mode may be useful, for example, in a system with complex configuration where static routes do not conflict with a firewall, while traffic flows following dynamic routes have to be subject to sophisticated firewall filtering.
- Auto : Both statically configured and dynamically learned Layer3 switching rules will be used. In this mode, maximum routing hardware acceleration is utilized.

Learn Mode

Synopsis: *string - one of the following keywords { host-oriented, flow-oriented }*

Default: *flow-oriented*

Defines how dynamically learned traffic flows are identified:

- flow-oriented: Traffic flows are identified by a 5-tuple signature:

```
Src IP address   +
Dst IP address   +
Protocol         +
Src TCP/UDP port +
Dst TCP/UDP port
```

  This mode should be used, if fine-granularity firewall filtering is configured in the device (i.e. some flows between two hosts should be forwarded, while other flows between the same two hosts should be filtered). However, this mode utilizes more Layer3 switching ASIC resources and is not recommended if fine-granularity firewall filtering is not required.

- host-oriented: Traffic flows are identified by a 2-tuple signature:

```
Src IP address   +
Dst IP address
```

  All traffic between two IP hosts is hardware-accelerated regardless of the protocol and TCP/UDP ports. This mode potentially controls multiple flows with a single rule and hence is more efficient in utilizing Layer3 switching ASIC resources.

Aging Time (sec)

Synopsis: *integer*

Default: *32*

This parameter configures the time a dynamically learned rule for a traffic flow, which has become inactive, is held before being removed from the Layer3 Switch forwarding table.

Static unicast routing is always enabled, while multicast routing is disabled by default. To change the Multicast Mode field to a value other than "disabled", you must first enable the Static Multicast Routing service. If the Static Multicast Routing service is not enabled, the system automatically overrides the Multicast Mode setting and changes it from "enabled" to "disabled".

32.2.2. Creating Static ARP Table Entries

To create static ARP table entries:

- In edit mode, navigate to */switch/layer3-switching/arp-table* and click **<Add arp-table>**.
- On the *Key settings* form, enter the network device IP address and click **Add**.
- On the *ARP Table Configuration* form, set the ARP table entry parameters.
- Commit the changes.

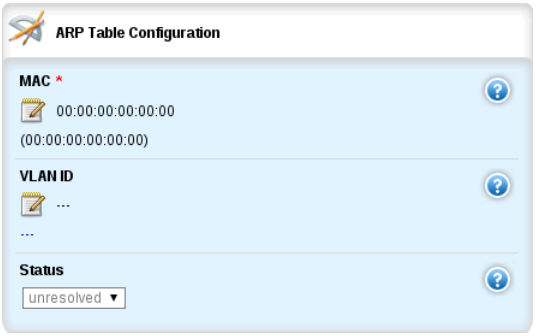

Figure 32.8. ARP Table Configuration form

MAC

Synopsis: *Unicast Ethernet MAC address in colon-separated hexadecimal notation*

Default: *00:00:00:00:00:00*

MAC address of the network device specified by the IP address.

VLAN ID

Synopsis: *integer*

VLAN Identifier of the VLAN upon which the MAC address operates.

status

Synopsis: *string - one of the following keywords { unresolved, resolved }*

Default: *unresolved*

ARP entry resolution status:

- Resolved : MAC-IP address pair is resolved and operational.
- Unresolved : the device hasn't resolved the MAC-IP address pair and keeps sending ARP requests periodically.

32.2.3. Viewing Static and Dynamic ARP Table Entries

The *ARP Table Summary* table lists all of the ARP table entries.

To view static and dynamic ARP table entries:

- Navigate to */switch/layer3-switching/arp-table-summary*.
- Review the entries on the *ARP Table Summary* form.

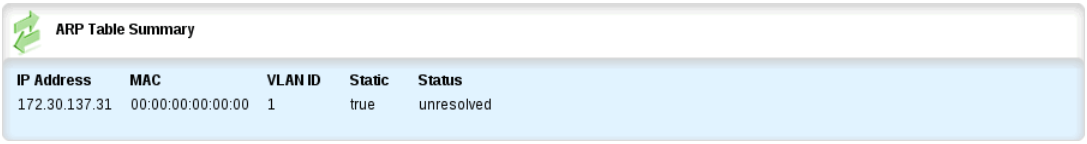

Figure 32.9. ARP Table Summary form

32.2.4. Viewing Routing Rules

The *Routing Rules Summary* table provides an overview of what is accelerated in the hardware; it shows what is actually going on inside the switch. The *Routing Rules Summary* form shows the details for a select rule.

To view the Routing Rules Summary table:

- Navigate to *switch/layer3-switching/routing-rules-summary*.
- Review the entries in the *Routing Rules Summary* table.

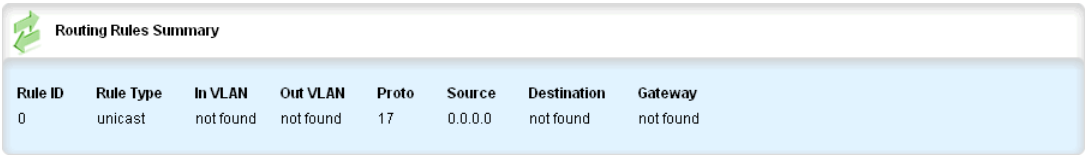

Figure 32.10. Routing Rules Summary table

To view the details for a routing rule:

- Navigate to *switch/layer3-switching/routing-rules-summary/{rule id}*.
- Review the entries on the *Routing Rules Summary* form.

Figure 32.11. Routing Rules Summary form

Rule Type

Synopsis: *string - one of the following keywords { hidden, invalid, unicast, multicast }*

Identifies the type of the rule: unicast,multicast,invalid

In VLAN

Synopsis: *integer*

Identifies the ingress VLAN. To match the rule, the packet's ingress VLAN must match the number.

Out VLAN

Synopsis: *integer*

Identifies the egress VLAN. The matched multicast packet is sent to the identified VLAN.

Proto

Synopsis: *unsigned byte*

IP Encapsulated Protocol number. Unless zero is specified, the incoming packet's IP protocol must match this number.

source

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *A string conforming to: "(([0-1]?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])\.){3}([0-1]?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])/\p{N}+"*

Synopsis: *string - the keyword { any }*

Identifies the source IP address or subnet. To match the rule, the incoming packet's source IP address must belong to the subnet.

destination

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *A string conforming to: "(([0-1]?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])\.){3}([0-1]?[0-9]?[0-9]|2[0-4][0-9]|25[0-5])/\p{N}+"*

Synopsis: *string - the keyword { any }*

Defines the destination IP address or subnet. To match the rule, the incoming packet's destination IP address must belong to the subnet.

gateway

Synopsis: *IPv4 address in dotted-decimal notation*

Defines the nexthop IP address. The matched unicast packet is sent to the identified gateway.

The address of the next-hop. The matched unicast packet is sent to the out-vlan. This is the list of VLANs the matched multicast packet will be sent to.

Pkts/sec

Synopsis: *unsigned integer*

Displays the statistical throughput of all packets matching the rule, in packets per second.

static

Synopsis: *boolean*

Whether the rule is static or dynamic.Static rules are configured as a result of management activity. Dynamic rules are automatically learned by the device and can be unlearned subject to Aging Time

routing-action

Synopsis: *string - one of the following keywords { exclude, forward }*

Action applied to packets matching the rule:

- Forward : perform hardware acceleration.
- Exclude : exclude from hardware acceleration and always pass matching packets to the CPU for software routing.

status

Synopsis: *string - one of the following keywords { excluding, pending, resolving, active }*

Whether the rule is currently operational or not:

- Active : rule is fully operational and can be applied, so hardware acceleration is performed.
- Resolving : rule is not operational yet due to some unresolved information, like ARP or gateway's MAC address in the MAC Address Table. Hardware acceleration is not performed.
- Pending : there are not enough hardware resources to setup the rule and all its dependencies. Hardware acceleration is not performed.

Insufficient - there are not enough hardware resources to set up the rule and all its dependencies. Hardware routing is not performed.

32.2.5. Flushing Dynamic Hardware Routing Rules

Flushing dynamic hardware routing rules removed dynamic rules from the Routing Rules Summary table.

> *You can only flush dynamic rules. Static rules, enabled by activating hardware acceleration, never age out. For more information on how to enable hardware acceleration, see Section 32.1, "Layer 3 Switching Fundamentals" .*

To flush dynamic hardware routing rules:

- Navigate to *switch/layer3-switching* and click **flush-dynamic-rules**.
- On the *Flush Dynamic Hardware Routing Rules* form, click **Perform**.

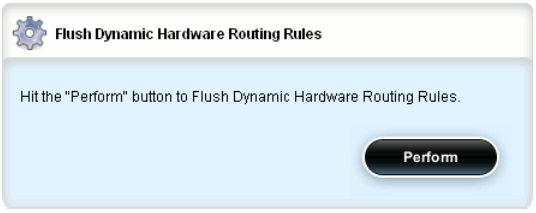

Figure 32.12. Flush Dynamic Hardware Routing Rules form

33. Tunnelling

Figure 33.1. Tunnelling menu

The tunnelling menu is accessible from the main menu under *tunnel*. This menu provides access to IPsec, L2TP, L2tunneld and GRE functions.

33.1. IPsec

33.1.1. VPN Fundamentals

IPsec (Internet Protocol SECurity) uses strong cryptography to provide both authentication and encryption services. Authentication ensures that packets are from the right sender and have not been altered in transit. Encryption prevents unauthorized reading of packet contents.

These services allow you to build *secure tunnels through untrusted networks.* Everything passing through the untrusted network is encrypted by the IPsec gateway and decrypted by the gateway at the other end. The result is a Virtual Private Network (VPN), a network which is effectively private even though it includes machines at several different sites connected by the insecure Internet.

The IPsec protocols were developed by the Internet Engineering Task Force (IETF) and are required as part of IP version 6.

Openswan is the open source implementation of IPsec used by ROX™.

The protocols used by IPsec are the Encapsulating Security Payload (ESP) and Internet Key Exchange (IKE) protocols.

ESP provides encryption and authentication (ensuring that a message originated from the expected sender and has not been altered on route).

IKE negotiates connection parameters, including keys, for ESP. IKE is based on the Diffie-Hellman key exchange protocol, which allows two parties without any initial shared secret to create one in a manner immune to eavesdropping.

33.1.1.1. IPsec Modes

IPSec has two basic modes of operation. In transport mode, IPSec headers are added as the original IP datagram is created. The resultant packet is composed of an IP header, IPSec headers and IP payload (including a transport header). Transport mode is most commonly used between IPsec end-stations, or between an end-station and a gateway.

In tunnel mode, the original IP datagram is created normally and then encapsulated into a new IP datagram. The resultant packet is composed of an new IP header, IPSec headers, old IP header and IP payload. Tunnel mode is most commonly used between gateways, the gateway acting as a proxy for the hosts behind it.

33.1.1.2. Policy-Based VPNs

RuggedBackbone™ supports the creation of policy-based VPNs, which may be characterized as follows:

- IPsec network interfaces are not created.
- The routing table is not involved in directing packets to the IPsec later.
- Only data traffic matching the tunnel's local and remote subnets is forwarded to the tunnel. Normal traffic is routed by one set of firewall rules and VPN traffic is routed based on separate rules.
- The firewall is configured with a VPN zone of type "IPsec".
- As IPsec packets are received, they are decoded, policy-flagged as IPsec-encoded, and presented as having arrived directly via the same network interface on which they were originally received.
- Firewall rules must be written to allow traffic to and from VPN tunnels. These are based on the normal form of source/destination IP addresses and IP protocol and port numbers. These rules, by virtue of the zones they match, use the policy flagging inserted by netkey and route matching data traffic to the proper interface.

33.1.1.3. Supported Encryption Protocols

Openswan supports the following standard encryption protocols:

- 3DES (Triple DES) – Uses three DES encryptions on a single data block, with at least two different keys, to get higher security than is available from a single DES pass. 3DES is the most CPU intensive cipher.
- AES – The Advanced Encryption Standard protocol cipher uses a 128-bit block and 128, 192 or 256-bit keys. This is the most secure protocol in use today, and is much preferred to 3DES due to its efficiency.

33.1.1.4. Public Key And Pre-shared Keys

In *public key cryptography*, keys are created in matched pairs (called public and private keys). The public key is made public while the private key is kept secret. Messages can then be sent by anyone who knows the public key to the holder of the private key. Only the owner of the private key can decrypt the message.

When you want to use this form of encryption, each router configures its VPN connection to use the RSA algorithm and includes the public signature of its peer. The RuggedBackbone™'s public signature is available from the output of the *tunnel/ipsec/display-public-key* menu.

In *secret key cryptography*, a single key known to both parties is used for both encryption and decryption.

When you want to use this form of encryption, each router configures its VPN connection to use a secret pre-shared key. The pre-shared key is configured through the *tunnel/ipsec/preshared-key* menu.

33.1.1.5. X509 Certificates

When one side of the VPN connection is placed from a dynamic IP (the so-called "roaming client"), X509 Certificates may be used to authenticate the connection. Certificates are digital signatures that are produced by a trusted source, namely a Certificate Authority (CA). For each host, the CA creates an certificate that contains CA and host information and "signs" the certificate by creating a digest of all the fields in the certificate and encrypting the hash value with its *private key*. The encrypted digest is called a "digital signature". The host's certificate and the CA *public key* are installed on all gateways that the host connects to.

When the gateway receives a connection request, it uses the CA *public key* to decrypt the signature back into the digest. It then recomputes its own digest from the plain text in the certificate and compares

the two. If both digests match, the integrity of the certificate is verified (it was not tampered with), and the public key in the certificate is assumed to be the valid public key of the connecting host.

33.1.1.6. NAT Traversal

Historically, IPSec has presented problems when connections must traverse a firewall providing Network Address Translation (NAT). The Internet Key Exchange (IKE) used in IPSec is not NAT-translatable. When IPSec connections must traverse a firewall IKE messages and IPSec-protected packets must be encapsulated as User Datagram Protocol (UDP) messages. The encapsulation allows the original untranslated packet to be examined by IPSec.

33.1.1.7. Other Configuration Supporting IPSec

If the router is to support a remote IPSec client and the client will be assigned an address in a subnet of a local interface, you must activate proxy ARP for that interface. This will cause the router to respond to ARP requests on behalf of the client and direct traffic to it over its connection.

IPSec relies upon the following protocols and ports:

- protocol 51, IPSEC-AH Authentication Header (RFC2402),
- protocol 50, IPSEC-ESP Encapsulating Security Payload (RFC2046),
- UDP port 500.

You must configure the firewall to accept connections on these ports and protocols. See Section 38.4, “Configuring The Firewall And VPN” in Chapter 38, *Firewall* for details.

33.1.1.8. The Openswan Configuration Process

Each VPN connection has two ends: the local router and the remote router. The Openswan configuration record describing a VPN connection can be used without change at either end. One side of the connection (typically the local side) is designated the “left” side and the other is designated the “right” side.

A convenient method is to configure both ends simultaneously with two command-line interface sessions (or two web browsers) open at the same time. The relevant information is the same in both sessions.

33.1.1.9. IPsec and Router Interfaces

If IPsec works on an interface which could disappear, such as a ppp connection, or if the IP address could change, you need to set the monitor-interface option for the IPsec connection. While this this option is set, IPsec will be restarted when the interface disappears and reappears or the IP address is changed.

For information on setting the monitor-interface option, see the Connection form at *tunnel/ipsec/connection/{line module}*.

33.1.1.10. L2TPD

L2TP stands for “Layer Two Tunneling Protocol”. The main purpose of this protocol is to tunnel PPP packets through an IP network, although it is also able to tunnel other layer 2 protocols.

On RuggedBackbone™, L2TPd is used in conjunction with Openswan and PPP to provide support for establishing a secure, private connection with the router using the Microsoft Windows VPN/L2TP client.

L2TPD listens on UDP port 1701. The firewall will need to be configured to allow connections to L2TPD via IPSec but to prevent connections to L2TPD directly without using IPsec.

33.1.2. IPsec Configuration

Figure 33.2. IPsec menu

The IPsec menu is accessible from the main menu under *tunnel*. The path to this menu is *tunnel/ipsec*. The IPsec form appears on the same screen as the *IPsec* menu.

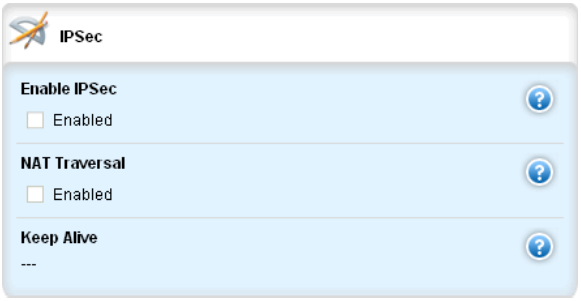

Figure 33.3. IPsec form

The IPsec form is used in configuring IPSec VPN.

Enable IPSec

Enables IPSec.

NAT Traversal

Enables NAT Traversal.

Keep Alive

Synopsis: *unsigned integer*

The delay (in seconds) for sending keepalive packets to prevent NAT router from closing its port when there is not enough traffic on the IPsec connection.

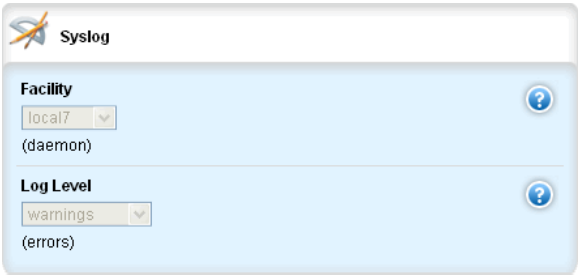

Figure 33.4. Syslog form

The path to the Syslog form is *tunnel/ipsec*.

Facility

Synopsis: *string - one of the following keywords { local7, local6, local5, local4, local3, local2, local1, local0, uucp, user, syslog, news, mark, mail, lpr, kern, daemon, cron, authpriv, auth }*

Default: *daemon*

The log facility.

Log Level

Synopsis: *string - one of the following keywords { warnings, notifications, informational, errors, emergencies, debugging, critical, alerts }*

Default: *errors*

The log level.

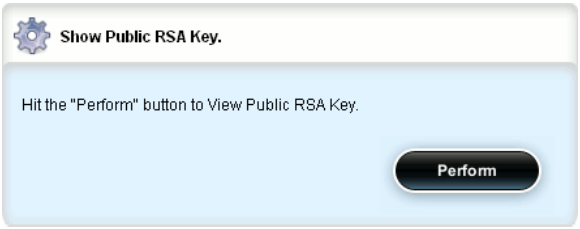

Figure 33.5. Show Public RSA Key form

The path to the Show Public RSA Key form is *tunnel/ipsec/display-public-key*. Clicking on the *Perform* button will display the public key.

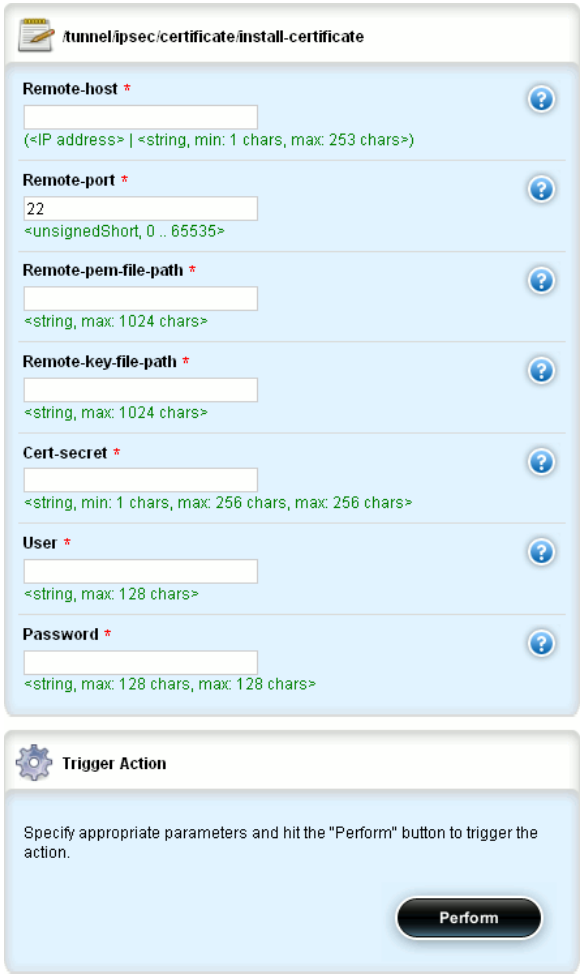

Figure 33.6. Install-Certificate forms

The path to the Install-Certificates forms is *tunnel/ipsec/certificate/install-certificate*. To install the certificates, enter the parameters and then click the *Perform* button.

Figure 33.7. Install-Ca-Certificate forms

The path to the Install-Ca-Certificate forms is *tunnel/ipsec/certificate/install-ca-certificate*. Enter the parameters and then click on the *Perform* button to install the certificates.

Figure 33.8. Install-Crl-File forms

The path to the Install-Crl-File forms is *tunnel/ipsec/certificate/install-crl-file*. To install the files, enter the parameters and then click the *Perform* button.

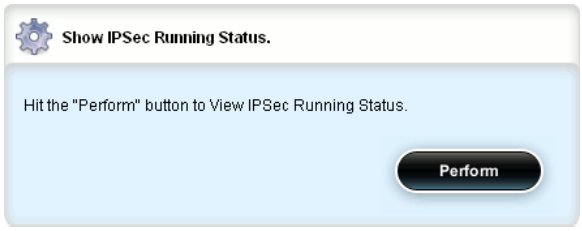

Figure 33.9. Show IPsec Running Status form

The path to the Show IPsec Running Status form is *tunnel/ipsec/status*. To display the IPsec status, click the *Perform* button.

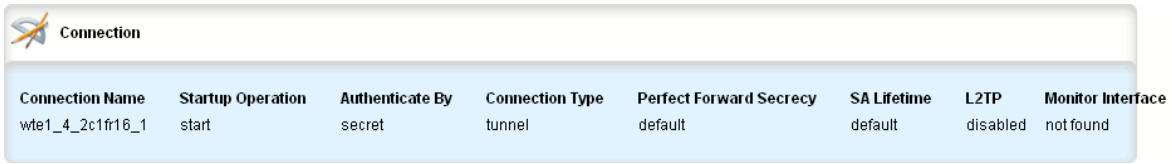

Figure 33.10. Connection table

If data is configured, the path to the Connection table will be *tunnel/ipsec/connection*.

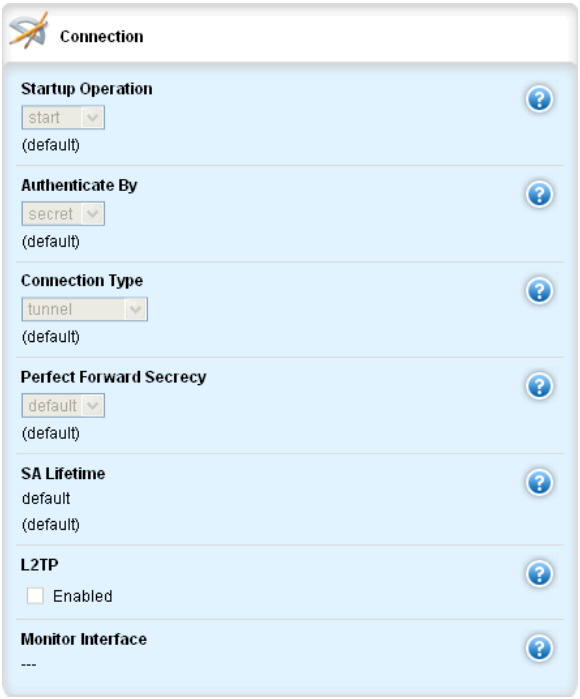

Figure 33.11. Connection form

If data is configured, the path to the Connection form will be *tunnel/ipsec/connection/{line module}*.

The Connection form is used for VPN connection configuration.

Connection Name

Synopsis: *string - the keyword { default }*

Synopsis: *A string conforming to: "[A-Za-z][A-Za-z0-9#%_\-+.,]+"*

The connection name. If the name is the default, this makes it the default setting for all connections.

Startup Operation

Synopsis: *string - one of the following keywords { default, route, start, add, ignore }*

Default: *default*

The action at IPSec startup time.

Authenticate By

Synopsis: *string - one of the following keywords { secret, rsasig, default }*

Default: *default*

The authentication method.

Connection Type

Synopsis: *string - one of the following keywords { default, passthrough, transport, tunnel }*

Default: *default*

The connection type. Tunnel signifies a host-to-host, host-to-subnet, or subnet-to-subnet tunnel; transport signifies a host-to-host transport mode; passthrough signifies that no IPsec processing should be done at all.

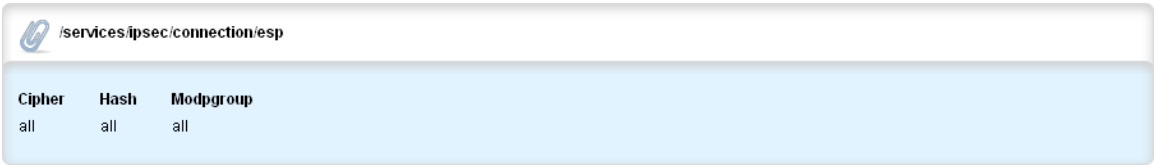

Figure 33.12. ESP table

If data is configured, the path to the ESP table will be *tunnel/ipsec/connection/{line module}/esp*.

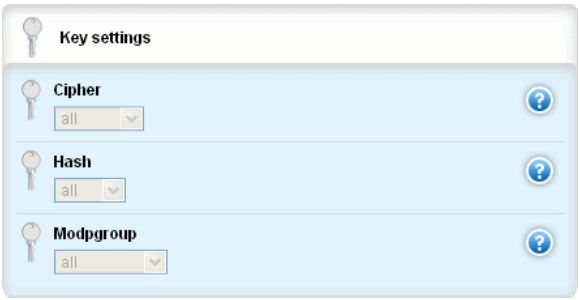

Figure 33.13. ESP Key Settings

If data is configured, the path to the ESP Key Settings form will be to click on *esp/{line module}*.

ESP pertains to the Phase 2 encryption/authentication algorithm to be used for the connection.

Cipher Algorithm

Synopsis: *string - one of the following keywords { any, aes128, aes192, aes256, aes, 3des }*

Cipher algorithm.

Hash Method

Synopsis: *string - one of the following keywords { any, md5, sha1 }*

Hash method.

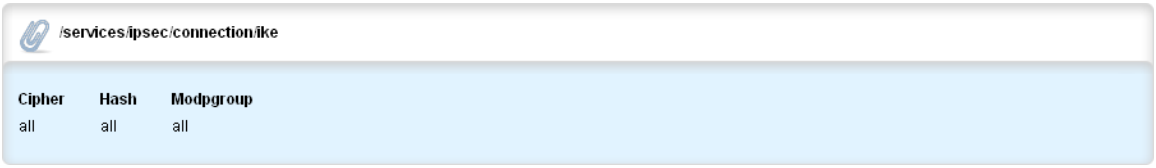

Figure 33.14. IKE table

If data is configured, the path to the IKE table will be *tunnel/ipsec/connection/{line module}/ike*.

IKE pertains to the Phase 1 encryption/authentication algorithm to be used for the connection. A Key Settings Form like the ESP Key Settings Form is also used for IKE.

Cipher Algorithm

Synopsis: *string - one of the following keywords { any, aes128, aes192, aes256, aes, 3des }*

Cipher algorithm.

Hash Method

Synopsis: *string - one of the following keywords { any, md5, sha1 }*

Hash method.

Modpgroup

Synopsis: *string - one of the following keywords { any, modp8192, modp6144, modp4096, modp3072, modp2048, modp1536, modp1024 }*

Modpgroup.

There are right side and left side IPsec forms. The forms for each side are used for IPSec system settings on each side. The forms are the same for both sides, so only the left side forms are shown here.

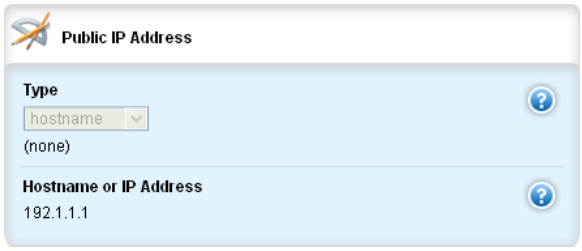

Figure 33.15. Public IP Address form

If data is configured, the path to the Public IP Address form will be *tunnel/ipsec/connection/{line module}/left*. The System Public Key, System Identifier and Nexthop to Other System forms appear on the same screen as the Public IP Address form.

The public-ip is the system identifier.

Type

Synopsis: *string - one of the following keywords { hostname, address, any, default-route, none }*

Default: *none*

Type.

Hostname or IP Address

Synopsis: *A string conforming to: "[^]+"*

Hostname or IP address.

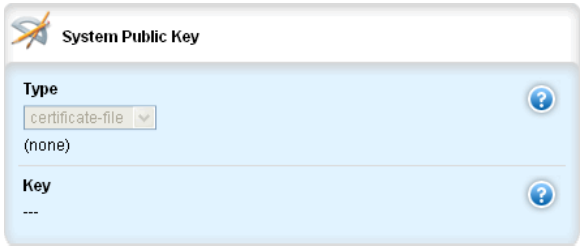

Figure 33.16. System Public Key form

It allows you to select the system's public key.

Type

Synopsis: *string - one of the following keywords { certificate-file, certificate-any, rsasig, none }*

Default: *none*

Type.

Key

Synopsis: *A string conforming to: "[^]+"*

Extra information.

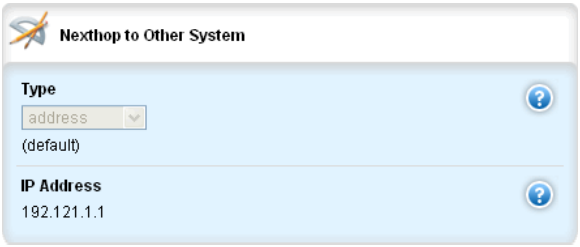

Figure 33.17. Nexthop To Other System form

Type

Synopsis: *string - one of the following keywords { address, default-route, default }*

Default: *default*

Type.

IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

IP address.

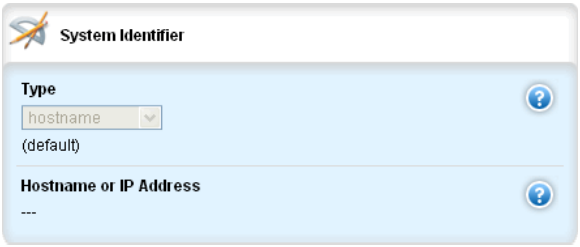

Figure 33.18. System Identifier form

type

Synopsis: *string - one of the following keywords { hostname, address, from-certificate, none, default }*

Default: *default*

Type.

Hostname or IP Address

Synopsis: *A string conforming to: "[^]+"*

Hostname or IP address.

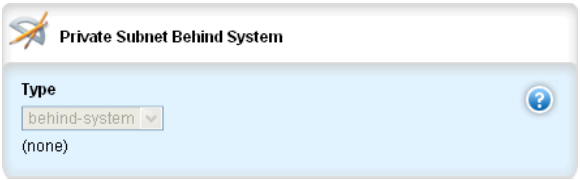

Figure 33.19. Private Subnet Behind System form

If data is configured, the path to the Private Subnet Behind System form will be *tunnel/ipsec/connection/{line module}/left/subnet*.

This form displays the private subnet behind the system.

Type

Synopsis: *string - one of the following keywords { behind-system, none }*

Default: *none*

Type.

Figure 33.20. Network table

The Network table displays a list of subnet addresses.

If data is configured, the path to the Preshared Key table will be *tunnel/ipsec/preshared-key*.

Figure 33.21. Preshared Key table

If data is configured, the path to the Preshared Key form will be *tunnel/ipsec/preshared-key/{line module}*.

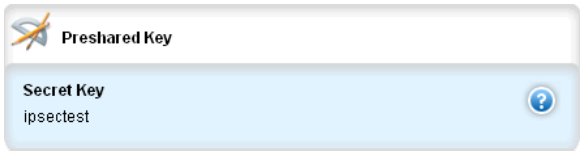

Figure 33.22. Preshared Key form

Remote Address

Synopsis: *string - the keyword { any }*

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *A string conforming to: "[^]+"*

The remote address.

Local Address

Synopsis: *string - the keyword { any }*

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *A string conforming to: "[^]+"*

The local address.

Secret Key

Synopsis: *"AES CFB128"-encrypted string*

The pre-shared key.

33.2. L2TP Tunnelling Configuration

Figure 33.23. L2TP menu

The path to the L2TP menu is *tunnel/l2tp*. The L2TP, DNS Server, PPP Options and WINS server forms appear on the same screen as this menu.

Figure 33.24. L2TP form

Enable L2TP

Enable L2TP.

Local IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

Local IP address.

First IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

First address in remote IP address pool.

Maximum Number of Connections

Synopsis: *unsigned integer*

Maximum number of connections.

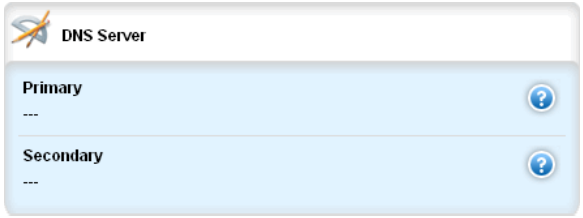

Figure 33.25. DNS Server form

Primary

Synopsis: *IPv4 address in dotted-decimal notation*

Primary DNS server.

Secondary

Synopsis: *IPv4 address in dotted-decimal notation*

Secondary DNS server.

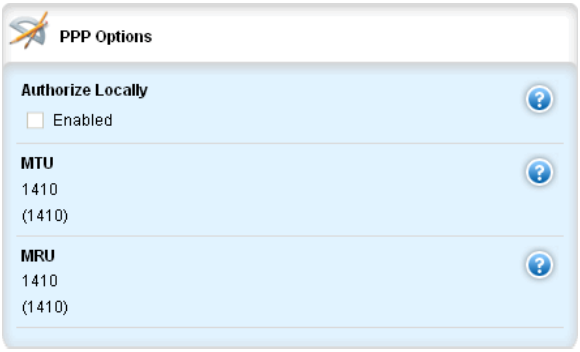

Figure 33.26. PPP Options form

Before enabling the Authorize Locally field on the PPP Options form, you need to add a PPP user name and password under the global/ppp/profiles/dialin menu. If you are not enabling the Authorize Locally field, you need to configure the Radius server for ppp authentication under the global/ppp/radius menu. For more information on PPP, see Chapter 13, PPP Users.

Authorize Locally

Authorize locally instead of using radius server.

MTU

Synopsis: *integer*

Default: *1410*

Maximum Transmit Unit (MTU) or maximum packet size transmitted.

MRU

Synopsis: *integer*

Default: *1410*

Maximum Receive Unit (MRU) or maximum packet size passed when received.

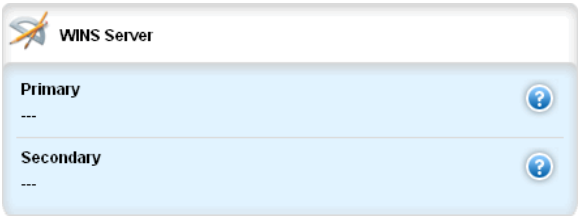

Figure 33.27. WINS Server form

Primary

Synopsis: *IPv4 address in dotted-decimal notation*

Primary WINS server.

Secondary

Synopsis: *IPv4 address in dotted-decimal notation*

Secondary WINS server.

33.3. Layer 2 Tunnelling

RuggedBackbone™ is capable of extending the range of services that communicate solely via Layer 2 protocols (i.e. at the level of Ethernet) by tunneling them over routed IP networks. The Layer 2 Tunnel Daemon supports the IEC61850 GOOSE protocol as well as a generic mechanism for tunneling by Ethernet type.

You can configure GOOSE tunnels and generic Layer 2 tunnels and also view tunnel status and statistics.

33.3.1. IEC61850 GOOSE Fundamentals

IEC61850 is an international standard for substation automation. It is a part of the International Electrotechnical Commission's (IEC) Technical Committee 57 (TC57) architecture for electric power systems. An important feature of IEC61850 is the fast transfer of event data. Transfers of Generic Substation Events (GSEs) are accomplished through the GOOSE (Generic Object Oriented Substation Event) protocol.

IEC61850 uses Layer 2 multicast frames to distribute its messages and hence is incapable of operating outside of a switched Ethernet Network. The GOOSE tunnel feature provides a capability to bridge GOOSE frames over a WAN.

GOOSE tunnels provide the following features:

- GOOSE traffic is bridged over the WAN via UDP/IP.
- One GOOSE traffic source can be mapped to multiple remote router Ethernet interfaces in mesh fashion.
- To reduce bandwidth consumption, GOOSE daemons may be located at each of the "legs" and at the center of a star network. The centrally located daemon will accept GOOSE packets and re-distribute them.
- Statistics report availability of remote GOOSE daemons, packet counts and Round Trip Time (RTT) for each remote daemon.
- When Virtual Router Redundancy Protocol (VRRP) is employed, GOOSE transport is improved by sending redundant GOOSE packets from each VRRP gateway.
- You can enable GOOSE forwarding by configuring a generic Layer 2 tunnel. When configured, RuggedBackbone™ listens for GOOSE packets on one VLAN and forwards them to another VLAN.

33.3.1.1. GOOSE Tunnel Implementation Details

The GOOSE protocol is supported by the Layer 2 Tunnel Daemon. The daemon listens to configured Ethernet interfaces and to the network itself (i.e. for tunnel connections from other daemon instances) on a configurable UDP port.

The Media Access Control (MAC) destination address of frames received from Ethernet is inspected in order to determine which GOOSE group they are in. The frames are then encapsulated in network headers and forwarded (with MAC source and destination addresses intact) to the network as GOOSE packets.

IEC61850 recommends that the MAC destination address should be in the range 01:0c:cd:01:00:00 to 01:0c:cd:01:01:ff.

GOOSE Packets received from the network are stripped of their network headers and forwarded to Ethernet ports configured for the same multicast address. The forwarded frames contain the MAC source address or the originating device, and not that of the transmitting interface. The VLAN used will be that programmed locally for the interface and may differ from the original VLAN. The frame will be transmitted with the highest 802.1p priority level (p4).

Packets received from the network will also be forwarded to any other remote daemons included in the group.

To enable forwarding for GOOSE packets, configure a generic Layer 2 tunnel to listen for GOOSE packets on one VLAN and forward them to a second VLAN. To configure the generic Layer 2 tunnel for this operation, set the following for the tunnel:

- **Ethernet Interface**: select the VLAN on which the GOOSE packets orginate.
- **Ethernet Type**: set as 0x8868. 0x8868
- **Remote Daemon**: select the VLAN to which to forward the GOOSE packets.

33.3.2. Generic Layer 2 Tunnel Fundamentals

The Layer 2 Tunnel Daemon also supports a generic mode of operation based on the Ethernet type of Layer 2 data traffic seen by the router. Multiple tunnels may be configured, each one with:

- Ethernet type
- Tunnel ingress (Ethernet interface)
- Tunnel egress (either another locally connected Ethernet interface, or the remote IP address of another Layer 2 Tunnel daemon instance running on another RuggedBackbone™)

33.3.2.1. Generic Tunnel Implementation Details

For each tunnel configured, the daemon monitors the specified Ethernet interface for Ethernet (Layer 2) frames of the specified type. If the configured egress is another local Ethernet port, frames are simply forwarded on that port, unmodified.

If the configured tunnel egress is a remote IP address, the daemon encapsulates the frames and forwards them to that address, where a corresponding Layer 2 Tunnel Daemon must be configured to receive tunneled frames for local retransmission. Encapsulation headers are stripped in order that the retransmitted frames are identical to those received at the tunnel ingress.

Other notes:

- Source and destination Ethernet MAC addresses are preserved, whether they are forwarded locally or remotely.
- Packets received from the network will also be forwarded to any other remote daemons included in the group.
- The UDP port number for inter-daemon communication must be the same throughout the network
- Enabling Generic L2 Tunneling on an Ethernet interface does not interfere with other (Layer 3) networking configuration on that interface, e.g. firewall rules, IP routing, etc.

Avoid network configurations where the daemons can form a traffic loop. The simplest such configuration is a triangle network where each daemon forwards to two other routers. Frames arriving at one router will start cycling in clockwise and counterclockwise directions.

To avoid such “packet storms”, frames forwarded to the network are tagged with an initial time to live count. The count is decremented at each relay to the network and prevents the frame from being relayed indefinitely.

33.3.3. Layer 2 Tunnelling Configuration

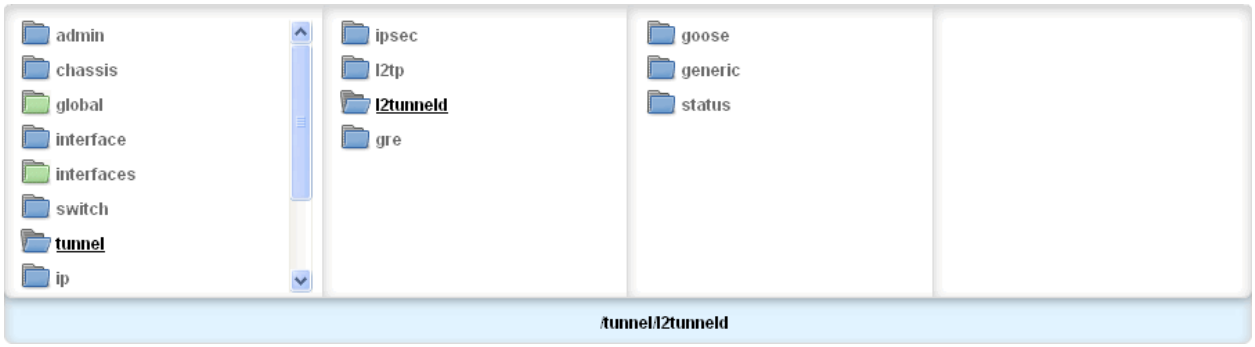

Figure 33.28. L2tunneld menu

The path to the L2tunneld (Layer 2) menu is *tunnel/l2tunneld*. The L2 Tunnel Daemon form appears on the same screen as this menu.

Figure 33.29. L2 Tunnel Daemon form

This form configures general settings for the daemon that apply to all supported tunnel configurations. The UDP-port field configures the port used by the daemon to communicate with other daemons. The Beacon-interval field configures how often a Round Trip Time (RTT) measurement message is sent to each remote peer. The interval takes the values “Off” to disable RTT measurement or a time of 10 – 3600 seconds.

All Layer 2 Tunnel daemons in the network must use the same port number. If the router employs a firewall, ensure that a hole is opened for each of the remote daemons using this port number.

enabled

Enables the L2 protocols server.

udp-port

Synopsis: *integer*

Default: *1311*

The UDP port to communicate with other deamon

beacon-interval

Synopsis: *string - the keyword { off }*

Synopsis: *integer*

Default: *60*

The Round Trip Time (RTT) of sending message

33.3.3.1. Goose

The forms and tables in this section are located under *tunnel/l2tunneld/goose*.

Figure 33.30. Goose Tunnel table

This table displays configured GOOSE tunnels.

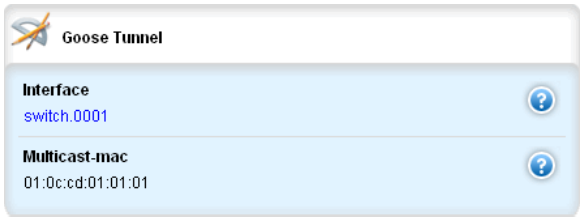

Figure 33.31. Goose Tunnel form

name

Synopsis: *A string*

Description of goose tunnel

interface

Synopsis: *A string*

The interface to listen on for goose frames

multicast-mac

Synopsis: *Multicast Ethernet MAC address in colon-separated hexadecimal notation*

The multicast MAC address to listen for

Figure 33.32. Remote Daemon of Goose Tunnel table

ip-address

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of remote L2 protocol server

33.3.3.2. Generic

The forms and tables in this section are located under *tunnel/l2tunneld/generic*.

Figure 33.33. Generic L2 Tunnel table

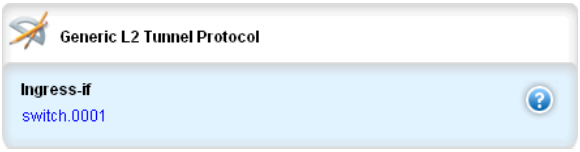

Figure 33.34. Generic L2 Tunnel Protocol form

name

Synopsis: *A string*

Description of goose tunnel

ingress-if

Synopsis: *A string*

The interface to listen on for Ethernet type frames

Figure 33.35. Generic L2 Tunnel Egress Interface table

egress-if

Synopsis: *A string*

Egress interface for Ethernet type frames

Figure 33.36. L2 Ethernet Type table

type

Synopsis: *string - the keyword { iso }*

Synopsis: *A string conforming to: "(0x[0-9A-Fa-f]{4})"*

Ethernet type to be forwarded (ie. 0xFEFE)

33.3.3.3. Status

The forms and tables in this section are located under *tunnel/l2tunneld/status/*. The submenus Goose, Generic and Round-Trip-Time are accessible from the *status* menu.

33.3.3.3.1. Status Goose

The forms and tables in this section are located under *tunnel/l2tunneld/status/goose*.

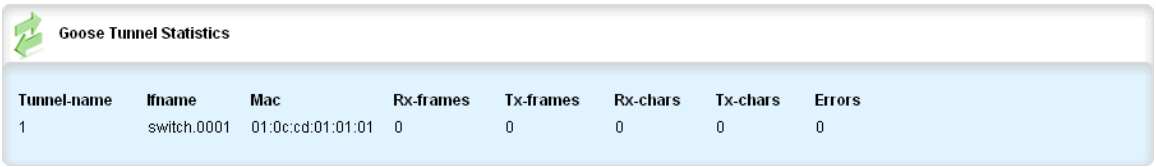

Figure 33.37. Goose Tunnel Statistics table

Figure 33.38. Goose Tunnel Statistics form

tunnel-name

Synopsis: *A string*

Goose Tunnel name

ifname

Synopsis: *A string*

VLAN Interface name

mac

Synopsis: *Multicast Ethernet MAC address in colon-separated hexadecimal notation*

Multicast Destination MAC Address of Goose message

rx-frames

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-frames

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-chars

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-chars

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

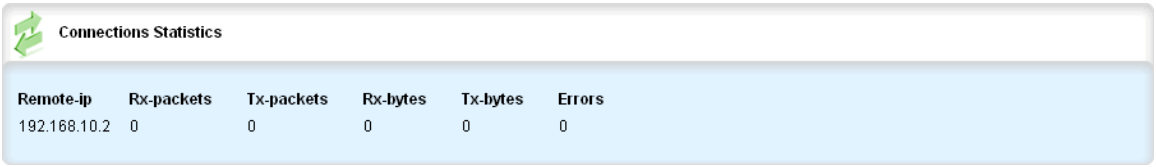

Figure 33.39. Connections Statistics table

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of remote goose daemon

rx-packets

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-packets

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-bytes

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-bytes

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

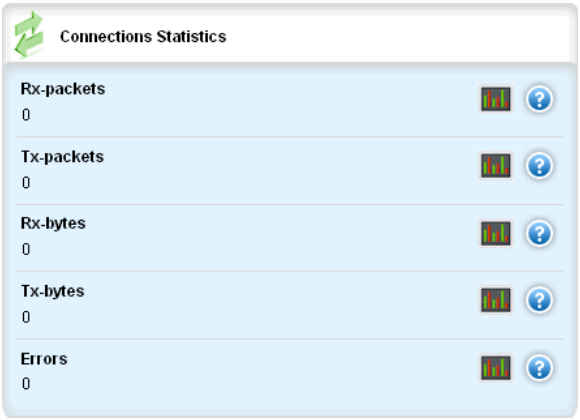

Figure 33.40. Connections Statistics form

33.3.3.3.2. Status Generic

The forms and tables in this section are located under *tunnel/l2tunneld/status/generic*.

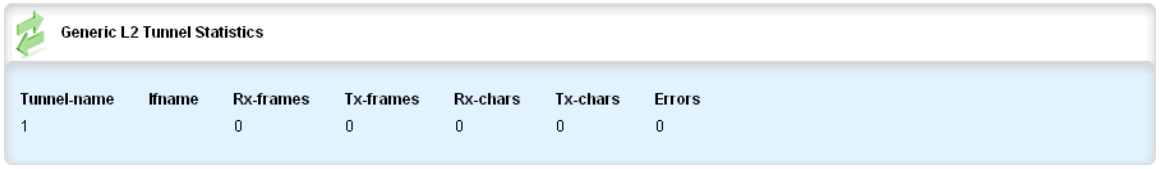

Figure 33.41. Generic L2 Tunnel Statistics table

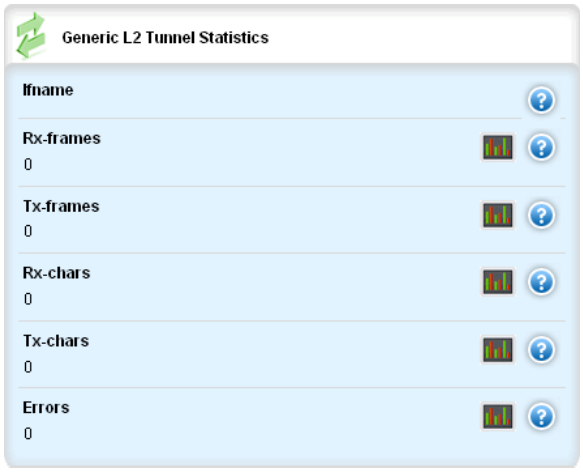

Figure 33.42. Generic L2 Tunnel Statistics form

tunnel-name

Synopsis: *A string*

Goose Tunnel name

ifname

Synopsis: *A string*

VLAN Interface name

rx-frames

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-frames

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-chars

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-chars

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

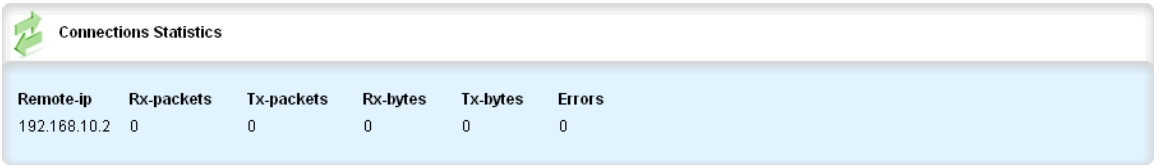

Figure 33.43. Connections Statistics table

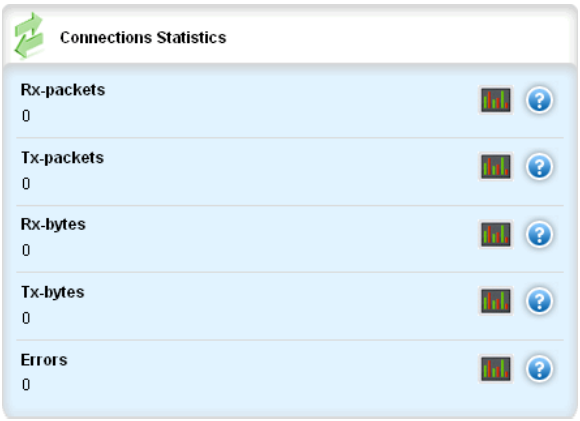

Figure 33.44. Connections Statistics form

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of remote goose daemon

rx-packets

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-packets

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-bytes

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-bytes

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

33.3.3.3.3. Status Round-Trip-Time

The forms and tables in this section are located under *tunnel/l2tunneld/status/round-trip-time*.

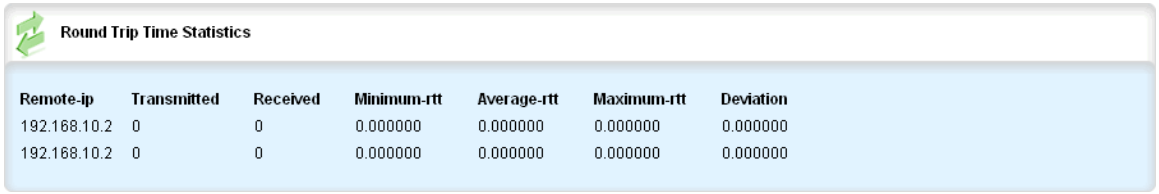
Round Trip Time Statistics
Remote-ip Transmitted Received Minimum-rtt Average-rtt Maximum-rtt Deviation
192.168.10.2 0 0 0.000000 0.000000 0.000000 0.000000
192.168.10.2 0 0 0.000000 0.000000 0.000000 0.000000

Figure 33.45. Round Trip Time Statistics table

The Round Trip Time Statistics table reflects the measured RTT to each remote daemon. The minimum, average, maximum and standard deviation of times is presented. Entries with a large difference between the Transmitted and Received fields indicate potential problems.

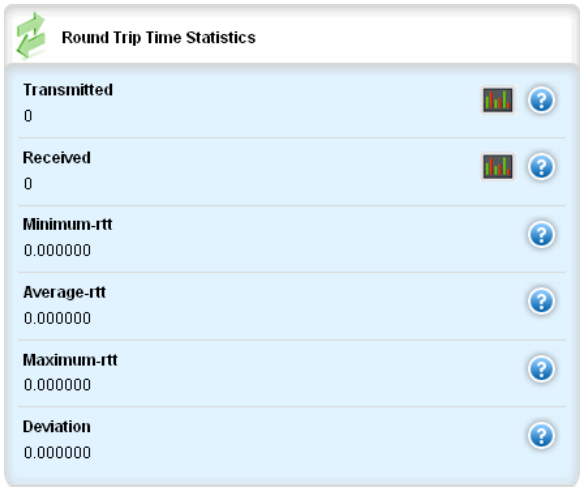

Figure 33.46. Round Trip Time Statistics form

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of remote goose daemon

transmitted

Synopsis: *unsigned integer*

The number of beacon frames transmitted over the tunnel

received

Synopsis: *unsigned integer*

The number of beacon frames transmitted over the tunnel

minimum-rtt

Synopsis: *A string*

The Minimum Beacon Round-Trip-Time

average-rtt

Synopsis: *A string*

The Average Beacon Round-Trip-Time

maximum-rtt

Synopsis: *A string*

The Maximum Beacon Round-Trip-Time

deviation

Synopsis: *A string*

The Standard Deviation

33.4. Generic Routing Encapsulation (GRE)

ROX™ is able to encapsulate multicast traffic and IPv6 packets and transport them through an IPv4 network tunnel.

A GRE tunnel can transport traffic through any number of intermediate networks. The key parameters for GRE in each router are the tunnel name, local router address, remote router address and remote subnet.

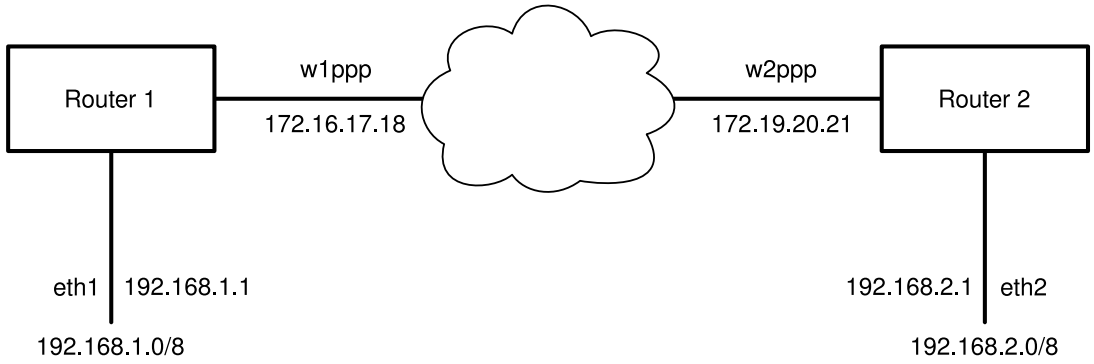

Figure 33.47. GRE Example

In the above example, Router 1 will use a GRE tunnel with a local router address of 172.16.17.18, a remote router address of 172.19.20.21, and a remote subnet of 192.168.2.0/24.

If you are connecting to a CISCO router (in place of Router 1 in the example above), the local router address corresponds to the CISCO IOS "source" address and the remote router address corresponds to the "destination" address.

You may also set a cost for the tunnel. If another method of routing between Router1 and Router2 becomes available, the tunnelled packets will flow through the lowest cost route. Optionally, you can restrict the packets by specifying the local egress device (in the case of router1, w1ppp).

33.4.1. Generic Routing Encapsulation Configuration

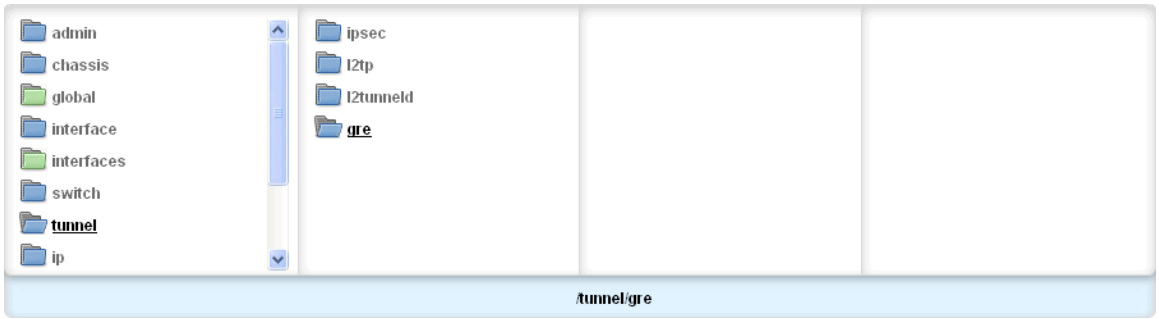

Figure 33.48. Generic Routing Encapsulation (GRE) menu

The path to this menu is *tunnel/gre*. If GRE tunnels are configured, they will be accessible via linked submenus.

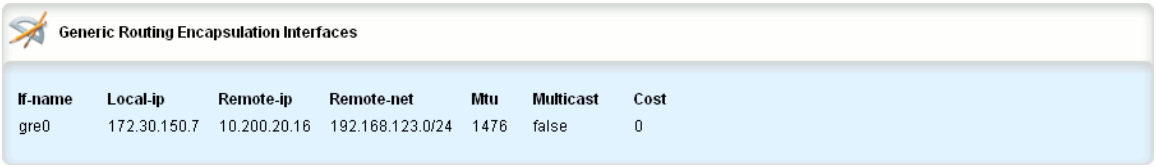

Figure 33.49. Generic Routing Encapsulation Interfaces table

The Generic Routing Encapsulation Interfaces table appears on the same screen as the *GRE* menu.

Figure 33.50. Generic Routing Encapsulation Interfaces form

The path to the Generic Routing Encapsulation Interfaces form is *tunnel/gre* and then clicking on one of the linked submenus that follow (for example, gre0).

if-name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

The GRE tunnel network interface name. The prefix 'gre-' will be added to this interface name.

local-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of the local end of the tunnel

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of the remote end of the tunnel

remote-net

Synopsis: *IPv4 address and prefix in CIDR notation*

The target network of the remote end of the tunnel (xxx.xxx.xxx.xxx/xx)

mtu

Synopsis: *integer*

Default: *1476*

The MTU of the GRE interface

multicast

Enables multicast traffic on the tunnel interface

cost

Synopsis: *integer*

Default:

The routing cost associated with networking routing that directs traffic through the tunnel

34. Dynamic Routing

34.1. Introduction

This chapter familiarizes the user with:

- Enabling the Dynamic Routing Suite
- Enabling and starting OSPF, RIP, and BGP
- Configuring OSPF, RIP, and BGP
- Obtaining OSPF, RIP, and BGP Status
- OSPF and VRRP

34.1.1. RIP, OSPF, and BGP

Dynamic routing is provided by a set of routing protocol daemons. The following daemons manage routing: core, ripd, ospfd, and bgpd.

The core daemon handles interfacing with the kernel to maintain the router's routing table and to check link statuses. It tells RIP, OSPF, and BGP what state links are in, what routes are in the routing table, and some information about the interfaces.

The ripd, ospfd, and bgpd daemons handle communications with other routers using the RIPv2, OSPFv2, and BGP protocols respectively, and decide which routers are preferred to forward to for each network route known to the router.

In complex legacy networks, RIP, OSPF, and BGP may be active on the same router at the same time. Typically, however, one of them is employed.

34.1.2. RIP Fundamentals

The Routing Information Protocol determines the best path for routing IP traffic over a TCP/IP network based on the number of hops between any two routers. It uses the shortest route available to a given network as the route to use for sending packets to that network.

The ROX™ RIP daemon (ripd) is an RFC1058 compliant implementation of RIP support RIP version 1 and 2. RIP version 1 is limited to obsolete class based networks, while RIP version 2 supports subnet masks as well as simple authentication for controlling which routers to accept route exchanges with.

RIP uses network and neighbor entries to control which routers it will exchange routes with. A network is either a subnet or a physical (broadcast-capable) network interface. Any router that is part of that subnet or connected to that interface may exchange routes. A neighbor is a specific router to exchange routes with specified by its IP address. For point to point links (T1/E1 links for example) one must use neighbor entries to add other routers to exchange routes with. The maximum number of hops between two points on a RIP network is 15, placing a limit on network size.

Link failures will eventually be noticed although it is not unusual for RIP to take many minutes for a dead route to disappear from the whole network. Large RIP networks could take over an hour to converge when link or route changes occur. For fast convergence and recovery, OSPF is a much better choice. RIP is a fairly old routing protocol and has mostly been superseded by OSPF.

34.1.3. OSPF Fundamentals

The Open Path Shortest First (OSPF) protocol routing determines the best path for routing IP traffic over a TCP/IP network based on link cost and quality. Unlike static routing, OSPF takes link failures and other network topology changes into account. Unlike the RIP routing protocol, OSPF provides less router to router update traffic.

The ROX™ OSPF daemon (ospfd) is an RFC 2178 compliant implementation of OSPFv2. The daemon also adheres to the RFC2370 (Opaque LSA) and RFC3509 (ABR-Types) extensions.

OSPF network design usually involves partitioning a network into a number of self-contained areas. The areas are chosen to minimize intra-area router traffic, making more manageable and reducing the number of advertised routes. Area numbers are assigned to each area. All routers in the area are known as Area routers. If traffic must flow between two areas a router with links in each area is selected to be an Area Border router, and serves as a gateway.

34.1.3.1. Link State Advertisements

When an OSPF configured router starts operating it issues a hello packet. Routers having the same OSPF Area, hello-interval and dead-interval timers will communicate with each other and are said to be neighbors

After discovering its neighbors, a router will exchange Link State Advertisements in order to determine the network topology.

Every 30 minutes (by default), the entire topology of the network must be sent to all routers in an area. If the link speeds are too low, the links too busy or there are too many routes, then some routes may fail to get re-announced and will be aged out.

Splitting the network into smaller areas to reduce the number of routes within an area or reducing the number of routes to be advertised may help to avoid this problem.

In shared access networks (i.e. routers connected by switches or hubs) a designated router and a backup designated are elected to receive route changes from subnets in the area. Once a designated router is picked, all routing state changes are sent to the designated router, which then sends the resulting changes to all the routers.

The election is decided based on the priority assigned to the interface of each router. The highest priority wins. If the priority is tied, the highest router-id wins.

34.1.4. Key OSPF And RIP Parameters

34.1.4.1. Network Areas

Network areas determine the regions within which routes are distributed to other routers. The subnets at a particular router can be added to its OSPF Area. The router will advertise these subnets to all routers in its area.

OSPF areas must be designed such that no single link failure will cause the network to be split into two disjoint networks.

A router can be part of multiple areas and function as a gateway between areas. When multiple areas are used on a network, area 0 is the backbone area. All areas must have a router connecting them to area 0.

34.1.4.2. Router-ID

Defines the ID of the router. By default this is the highest IP assigned to the router. It is often a good idea to configure this value manually to avoid the router-id changing if interfaces are added or deleted from the router. During elections for designated router, the router-id is one of the values used to pick the winner. Keeping the router-id fixed will avoid any unexpected changes in the election of the master router.

34.1.4.3. Hello Interval and Dead Interval

The hello interval is the time between transmission of OSPF Hello packets. The dead interval is the time to wait without seeing an OSPF Hello packet before declaring a neighboring router dead and discarding its routes. It is recommended that the dead interval be at least four times the hello interval for reliable operation.

Lower values of these settings will help to speed up the change in network routes when the topology of the network changes. It will also increase the load on the router and the links, due to higher traffic caused by the increase in messages. Lower values will also put limits on the number of routes that can be distributed within an area, as will running over slower links.

OSPF will not work properly if the Hello Interval and Dead Interval are not identical on every router in an area.

34.1.4.4. Active/Passive Interface Default

OSPF regards router interfaces as either passive or active, sending OSPF messages on active interfaces and ignoring passive interfaces. By default, newly created interfaces are viewed as passive from OSPF until they are configured active. This is more efficient and secure for the router. The default type for new interfaces is controlled by the passive interface default option in the OSPF Global Parameters.

The default setting of Passive Interface Default means that you must explicitly configure interfaces active before OSPF will attempt to use them.

34.1.4.5. Redistributing Routes

Routes for subnets which are directly connected to the router but are not part of the OSPF area or RIP or BGP networks can be advertised if "redistribute connected" is enabled in the OSPF, RIP, or BGP Global Parameters. Static routes and routes handled by the kernel can also be redistributed if "redistribute static" and "redistribute kernel" are enabled, respectively.

34.1.4.6. Link Detect

Link detection is enabled for active network interfaces. Link detection ensures that the appropriate routing daemon is notified when an interface goes down and stops advertising subnets associated with that interface. The routing daemon resumes advertising the subnet when the link is restored. This allows routing daemons to detect link failures more rapidly, as the router does not have to wait for a “dead interval” to time out). Link detection also causes ”redistributed“ routes to start and stop being advertised based on the status of their interface links.

34.1.4.7. Configuring OSPF Link Costs

Link cost determines which route to use when multiple links can reach a given destination. By default, OSPF assigns the same cost to all links unless it is provided with extra information about the links. Each interface is assumed to be 10Mbit unless specified otherwise in the auto-cost bandwidth setting found under the *ip* menu.

The default OSPF reference bandwidth for link cost calculations is 100Mbit. The reference bandwidth divided by the link bandwidth gives the default cost for a link,which by default is 10. If a specific bandwidth is assigned to each link, the costs take this into account.

You can manually assign a link cost in the OSPF Interface Configuration for each interface. Do this for cases where the speed of the link should not be used as the method for choosing the best link.

34.1.4.8. OSPF Authentication

OSPF authentication is used when it is desirable to prevent unauthorized routers from joining the OSPF network. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the OSPF network. Authentication adds a small overhead due to the encryption of messages, so is not to be preferred on completely private networks with controlled access.

34.1.4.9. RIP Authentication

RIP authentication is used when it is desirable to prevent unauthorized routers from joining the network. RIP authentication is supported by per-interface configuration or the use of key-chains. Separate key chains spanning different groups of interfaces and having separate lifespans are possible. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the RIP network.

34.1.4.10. Administrative Distances

The router may work with different routing protocols at the same time, as well as employing local interface and statically assigned routes. An administrative distance, (from 0 to 255) is a rating of the trustworthiness of a routing information source. For a given route, the protocol having the lowest administrative distance will be chosen. By default the distances for a connected interface is 0 and for a static route is 1. By default, OSPF will set an administrative distance of 110 and RIP will set a distance of 120.

34.1.5. OSPF And VRRP Example Network

This network consists of three routers connected in a ring with T1/E1 links. Router 1 and 2 and the switched network represent a remote site in which the routers supply a redundant gateway to the hosts via VRRP and the T1/E1 links supply a redundant network connection to the rest of the network.

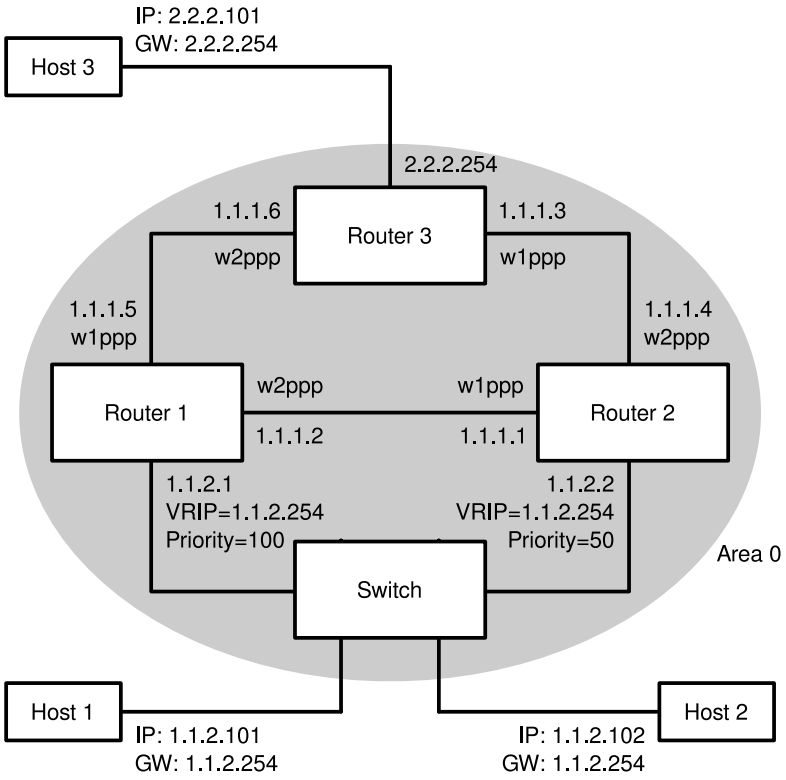

Figure 34.1. OSPF and VRRP Example

34.1.5.1. Area And Subnets

As the OSPF design is simple, an area of 0 is used. The three point-to-point T1/E1 links are placed in the area by adding 1.1.1.0/24 to it. Router 1 and 2 will include their Ethernet links by adding subnet 1.1.2.0/24 to their area descriptions. Router 3 must also include 2.2.2.0/24 in its area description so that its existence is advertised.

The point-to-point T1/E1 interfaces and Ethernet interfaces on Router 1 and 2 must be made active. The Ethernet interface on Router 3 can be left passive since it does not participate in OSPF advertisements.

Router 1 and 2 must enable link-detect, to stop advertising 1.1.1.0/24 in the event of a link failure.

34.1.5.2. VRRP Operation

Router 1 and 2 have VRRP setup on their Ethernet connection so that they can both function as the gateway for the clients on their network segment. Normally Router 1 is the VRRP master, and only in case of a link failure to the switch or the router failing, will Router 2 take over the virtual IP. The virtual IP used as the gateway is 1.1.2.254. Each router also has its own IP on the network so that each can be reached individually.

If Router 1 or its Ethernet link fail, VRRP will detect the link being down and remove the direct route to the 1.1.2.0/24. VRRP on Router 2 will stop seeing messages from Router 1, elect itself master and will take over the gateway for the network.

OSPF on router 1 will notice the link being down (and the route to 1.1.2.0/24 disappearing) and will use information from router 2 install a route to 1.1.2.0/24 via Router 2.

Router 3 will notice that Router 2 is now a more direct path to 1.1.2.0/24 network and start sending to Router 2 instead of Router 1.

After the failure all routers still know how to reach the entire network, and the clients on 1.1.2.0/24 can still send on the network using the same gateway address. The clients will see only a MAC address change of the gateway and experience a few seconds of network outage When the link returns, VRRP will switch back to the master, and the routes will return to their normal state.

Note that if the Router 1 WAN link fails, Router will see routes to Router3 via the Router 1 – Router 2 WAN and Ethernet links. If the faster Router 1 – Router 2 Ethernet path fails, Router 1 will fall back to the Router 1 – Router 2 WAN link.

Note that it would not be useful to leave the Ethernet 1.1.2.0/24 subnets out of the area and turn on redistribute connected as OSPF would not use the subnets for routing.

34.1.6. BGP Fundamentals

The Border Gateway Protocol (BGP, RFC 4271) is a robust and scalable routing protocol. BGP is designed to manage a routing table of up to 90000 routes. Therefore, it is used in large networks or among groups of networks which have common administrative and routing policies. If BGP is used to exchange routing information between different networks, it is called Exterior BGP (EBGP). Interior BGP (IBGP) is used to exchange routing information between routers within the same network.

34.2. Dynamic Routing Configuration

Figure 34.2. Dynamic Routing Menu

The Dynamic Routing menu is accessible from the main menu under *routing*. The path to this menu is *routing/dynamic*.

The *BGP*, *RIP* and *OSPF* menus provide access to features that enable configuration of the corresponding routing protocols.

34.3. RIP

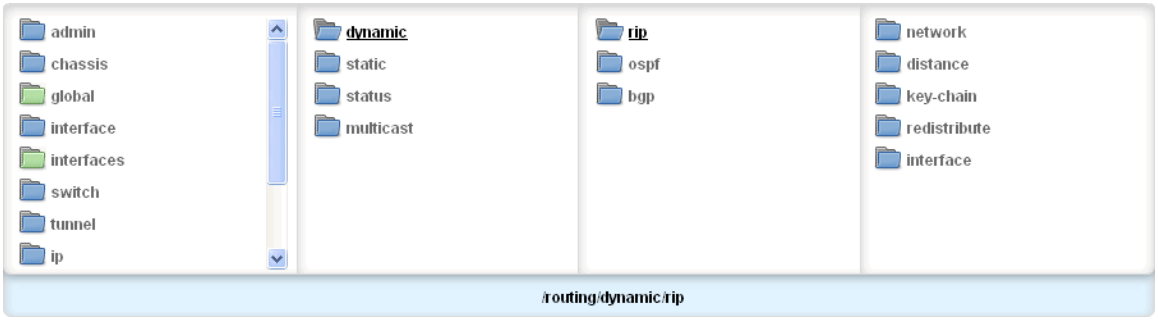

Figure 34.3. RIP Menu

The RIP menu is accessible from the main menu under *routing*. The path to this menu is *routing/dynamic/rip*.

34.3.1. RIP Configuration

The RIP Configuration form and Routing Timers form appear on the same screen as the *RIP* menu.

Figure 34.4. RIP Configuration Form

Enable RIP

Enables the RIP dynamic routing protocol.

Default Information Originate

The route element makes a static route only inside RIP. This element should be used only by advanced users who are particularly knowledgeable about the RIP protocol. In most cases, we recommend creating a static route and redistributing it in RIP using the redistribute element with static type.

Default Metric

Synopsis: *integer in the range [-32768 to 32767]*

Default: *1*

This element modifies the default metric value for redistributed routes. The value does not affect connected routes even if it is redistributed by redistribute connected. To modify the connected route's metric value, please use the redistribute connected metric.

Distance Default

Synopsis: *unsigned integer*

Set the default RIP distance to a specified value.

Version

Synopsis: *integer in the range [-32768 to 32767]*

Set the RIP version to accept for reads and send. The version can be either 1 or 2.

Disabling RIPv1 by specifying version 2 is STRONGLY encouraged.

Figure 34.5. Routing Timers Form

The RIP protocol has several timers. The user can configure those timers' values by the timers-basic element.

The default settings for the timers are as follows:

* The update timer is 30 seconds. Every update timer seconds, the RIP process is awakened to send an unsolicited response message containing the complete routing table to all neighboring RIP routers.

* The timeout timer is 180 seconds. Upon expiration of the timeout, the route is no longer valid; however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped.

* The garbage collect timer is 120 seconds. Upon expiration of the garbage-collection timer, the route is finally removed from the routing table.

Update Timer

Synopsis: *unsigned integer*

Default: *30*

The routing table update timer (in seconds).

Timeout Timer

Synopsis: *unsigned integer*

Default: *180*

The routing information timeout timer (in seconds).

Garbage Collection Timer

Synopsis: *unsigned integer*

Default: *120*

The garbage collection timer (in seconds).

34.3.1.1. RIP Networks

Neighbors are specific routers with which to exchange routes using the RIP protocol. This can be used when you want to explicitly control which routers are part of your RIP network.

Networks are used when you want to add any router that is part of a specific subnet, or connected to a specific network interface to be part of your RIP network.

Both neighbors and networks can be used at the same time.

For point to point links (T1/E1 links for example), one must use neighbor entries to add other routers to exchange routes with. Also note that RIP v1 does not send subnet mask information in its updates. Any defined networks are restricted to the classic (in the sense of Class A, B and C) networks. RIP v2 does not have this failing.

Subnet

Subnet Address/Prefix

Synopsis: *IPv4 address and prefix in CIDR notation*

Network address/prefix.

Neighbor

Neighbor IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

Neighbor IP address.

34.3.1.2. Distance

Distance with Matched Subnet

Subnet/Prefix

Synopsis: IPv4 address and prefix in CIDR notation

IP Address/Prefix.

Distance

Synopsis: unsigned integer

Distance value.

34.3.1.3. RIP Key Chains

The Key Chains table configures authentication keys used on the interfaces. By defining the keys in a key chain, the same settings can be applied to multiple groups of interfaces. Without key chains the same settings would have to be entered for each interface separately.

Key chains also allow multiple keys to be entered in a single key chain with a start time for when that key should become valid as well as the duration the key is valid. This allows multiple keys to be set up with automatic transitions from one key to the next over time.

A key consists of a key string, which is the value used for authentication. It also has the optional lifetime to accept RIP messages with the key, and the optional lifetime to send RIP messages with that key.

Key chain management.

Key Chain Name

Synopsis: *string*

The key chain name.

Key Configuration

Configure a key.

Key ID

Synopsis: *unsigned integer*

The key identifier number.

Key

Synopsis: *string*

Sets the key string.

Accept Lifetime

Set the accept lifetime of the key.

Time to Start

Synopsis: *date and time specification*

The time to start.

Expire Time

Synopsis: *date and time specification*

Synopsis: *string - the keyword { infinite }*

Expire time.

Send Lifetime

Set the send lifetime of the key.

Time to Start

Synopsis: *date and time specification*

The time to start.

Expire Time

Synopsis: *date and time specification*

Synopsis: *string - the keyword { infinite }*

Expire time.

34.3.1.4. Redistribute

This element redistributes routing information into the RIP tables from route entries specified by type.

Redistribute Route from other Protocols

Redistribute Type

Synopsis: *string - one of the following keywords { bgp, ospf, connected, static, kernel }*

Redistribute route type.

Metric

Synopsis: *integer in the range [-32768 to 32767]*

The metric for redistributed routes.

34.3.1.5. RIP Interfaces

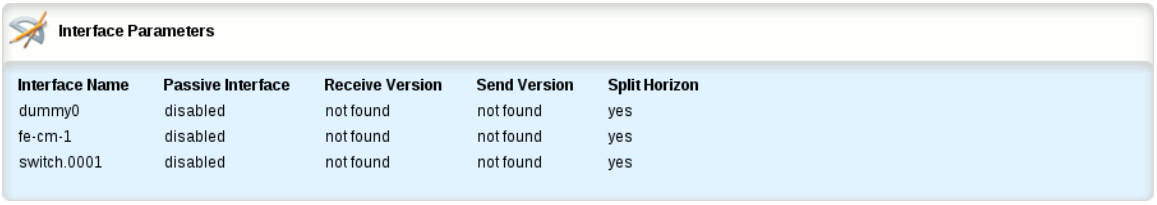

Interface Parameters

Interface Name	Passive Interface	Receive Version	Send Version	Split Horizon
dummy0	disabled	not found	not found	yes
fe-cm-1	disabled	not found	not found	yes
switch.0001	disabled	not found	not found	yes

Figure 34.6. RIP Interface Parameters Table

RIP parameters refer to RIP parameters on the selected interface. The path to the RIP Interface Parameters table is *routing/dynamic/rip/interface*.

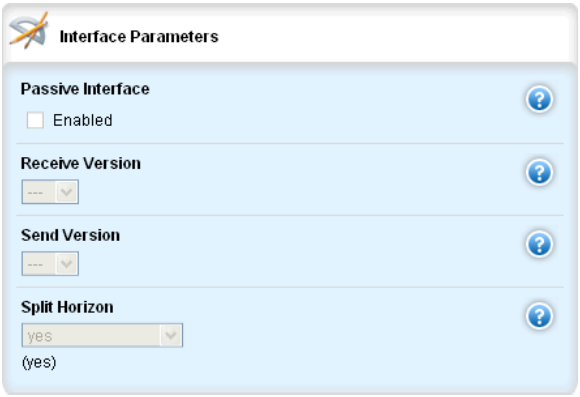

Figure 34.7. RIP Interface Parameters Form

To display the RIP Interface Parameters form and the Authentication form, navigate to *routing/dynamic/rip/interface/{interface}*.

Passive Interface

The specified interface is set to passive mode. In passive mode, all received packets are processed normally and ripd sends neither multicast nor unicast RIP packets except to RIP neighbors specified with a neighbor element.

Receive Version

Synopsis: *string - one of the following keywords { 2,1, 1,2, 2, 1 }*

The version of RIP packets that will be accepted on this interface. By default, version 1 and version 2 packet will be accepted.

Send Version

Synopsis: *string - one of the following keywords { 2,1, 1,2, 2, 1 }*

The version of RIP to send packets with. By default, version 2 packet will be sent.

Split Horizon

Synopsis: *string - one of the following keywords { poisoned-reverse, no, yes }*

Default: *yes*

A split horizon.

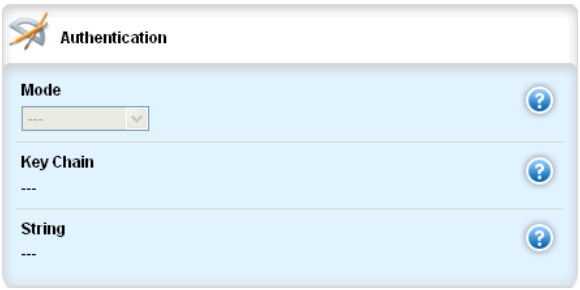

Figure 34.8. Authentication Form

Mode

Synopsis: *string - one of the following keywords { none, text, md5-old-ripd, md5-rfc }*

The authentication mode.

Key Chain

Synopsis: *string*

The authentication key-chain.

String

Synopsis: *string*

The authentication string.

> *Split horizon controls whether routes learned through an interface should be allowed to be advertised back out that interface. By default RIP advertises all routes it knows about to everyone, which makes it take a very long time for dropped links to age out of the network. The split horizon prevents advertising those routes back out the same interface which helps to control this problem. Some network topologies with rings of routers will still have some issues with aging out dead routes even with split horizon enabled but they will still age out faster. If fast network recovery is desired, use OSPF.*

34.4. OSPF

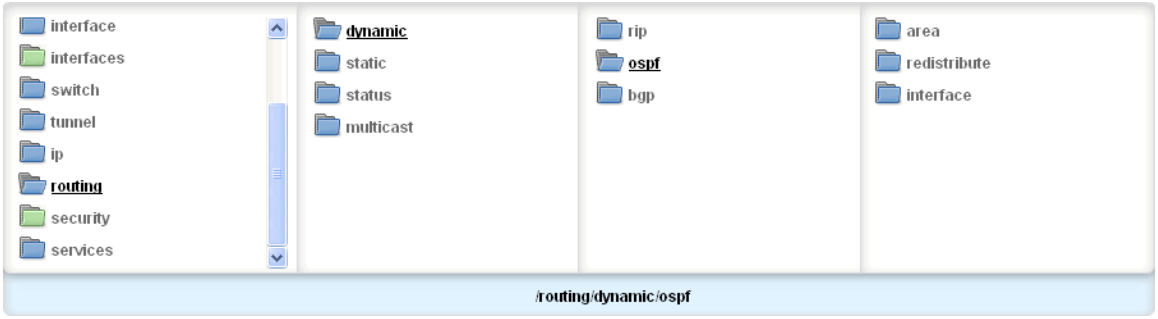

Figure 34.9. OSPF Menu

The OSPF menu is accessible from the main menu under *routing*. The path to this menu is *routing/dynamic/ospf*. The OSPF Area Distance form and the OSPF Configuration form appear on the same screen as the OSPF menu.

34.4.1. OSPF Configuration

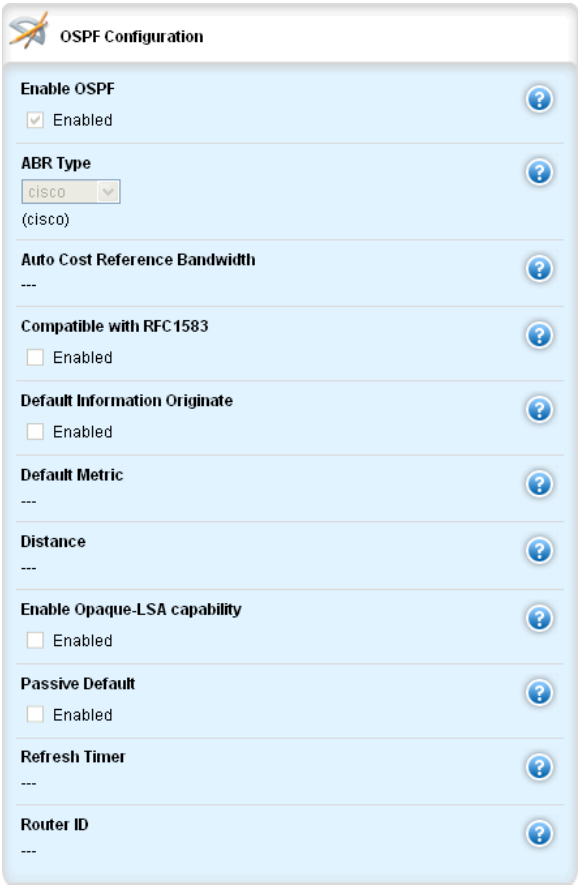

Figure 34.10. OSPF Configuration Form

Enable OSPF

Enables the OSPF dynamic routing protocol.

ABR Type

Synopsis: *string - one of the following keywords { standard, shortcut, ibm, cisco }*

Default: *cisco*

The OSPF ABR type.

Auto Cost Reference Bandwidth

Synopsis: *unsigned integer*

Default: *100*

Calculates the OSPF interface cost according to bandwidth [1-4294967 Mbps]

Compatible with RFC1583

Enables the compatibility with the obsolete RFC1583 OSPF (the current is RFC2178)

Default Information Originate

Advertises the default route.

Default Metric

Synopsis: *unsigned integer*

The default metric of redistribute routes.

Distance

Synopsis: *unsigned integer*

The administrative distance.

Enable Opaque-LSA capability

Enables the Opaque-LSA capability (rfc2370).

Passive Default

Suppresses routing updates on interfaces by default.

Refresh Timer

Synopsis: *unsigned short integer*

Default: *10*

The refresh timer.

Router ID

Synopsis: *IPv4 address in dotted-decimal notation*

The Router ID for OSPF.

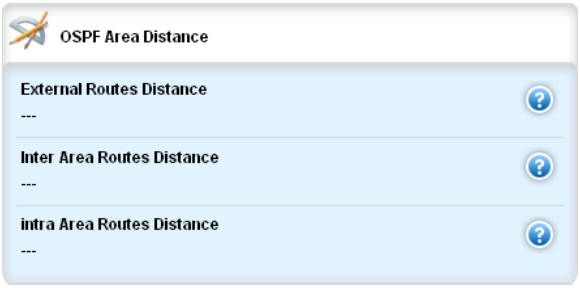

Figure 34.11. OSPF Area Distance Form

The OSPF Area Distance form can be used to define OSPF external, inter-area or intra-area routes distance.

External Routes Distance

Synopsis: *unsigned integer*

The administrative distance for external routes.

Inter Area Routes Distance

Synopsis: *unsigned integer*

The administrative distance for inter-area routes.

intra Area Routes Distance

Synopsis: *unsigned integer*

The administrative distance for intra-area routes.

34.4.1.1. OSPF Network Areas

OSPF uses areas to control which routes are distributed between routers.

Area ID.

Synopsis: *IPv4 address in dotted-decimal notation*

The OSPF Area ID (format: A.B.C.D).

Area Network/Prefix

Synopsis: *IPv4 address and prefix in CIDR notation*

The OSPF area network/prefix.

34.4.1.2. OSPF Redistribute

Redistribute from other Routing Protocols

This feature redistributes information from another routing protocol.

Redistribute Route From

Synopsis: *string - one of the following keywords { bgp, rip, connected, static, kernel }*

Redistributes the route type.

Metric Type

Synopsis: *integer in the range [-32768 to 32767]*

Default: *2*

The OSPF exterior metric type for redistributed routes.

Metric

Synopsis: *unsigned integer*

The metric for redistributed routes.

34.4.1.3. OSPF Interfaces

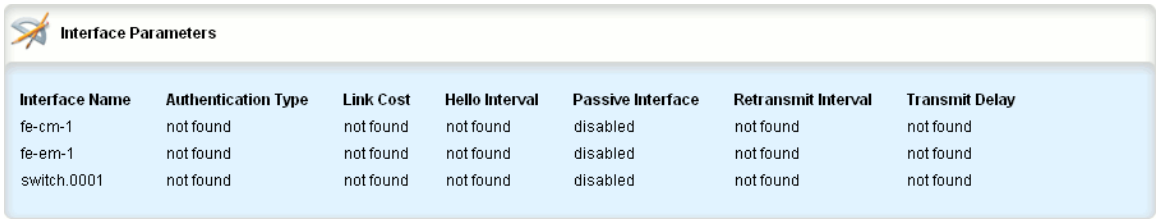

Figure 34.12. Interface Parameters Table

The path to the Interface Parameters table is *routing/dynamic/ospf/interface*.

Figure 34.13. Interface Parameters Form

The path to the Interfaces form and Dead Interval form is *routing/dynamic/ospf/interface/{interface}*.

OSPF parameters on the selected interface.

Interface Name

Synopsis: *string*

Interface name.

Authentication Type

Synopsis: *string - one of the following keywords { null, message-digest }*

The authentication type on this interface.

Link Cost

Synopsis: *unsigned integer*

The link cost. If not set, it cost is based on reference bandwidth.

Hello Interval

Synopsis: *unsigned integer*

Default: *10*

The time (in seconds) between sending hello packets.

priority

Synopsis: *unsigned byte integer*

Default: *1*

Priority of interface.

Passive Interface

Whether an interface is active or passive. Passive interfaces do not send LSAs to other routers and are not part of an OSPF area.

Retransmit Interval

Synopsis: *unsigned integer*

Default: *5*

Time (in seconds) between retransmitting lost link state advertisements.

Transmit Delay

Synopsis: *unsigned integer*

Default: *1*

The link state transmit delay (in seconds).

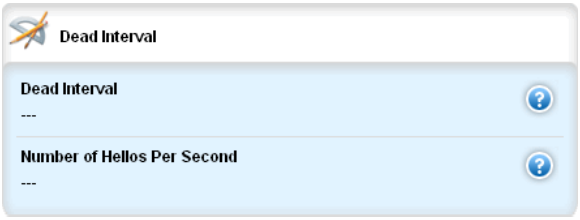

Figure 34.14. Dead Interval Form

A dead interval is the time before considering a router dead.

Dead Interval

Synopsis: *unsigned integer*

Default: *40*

The time before considering a router dead (in seconds).

Number of Hellos Per Second

Synopsis: *unsigned byte integer*

The number of times a hello message can be sent within one second.

Configuration Parameters

Configuration forms and display tables can be found at *routing/dynamic/ospf/interface*, then clicking on one of the interface submenus (for example, dummy0) and then clicking on the further set of submenus that follow (authentication-ip, cost-ip, dead-interval-ip, hello-interval-ip, message-digest-key, message-digest-key-ip, retransmit-interval-ip and transmit-delay-ip).

34.5. BGP

34.5.1. BGP configuration

Figure 34.15. BGP Menu

To display the BGP menu, navigate to *routing/dynamic/bgp*.

Figure 34.16. BGP Configuration Form

The BGP Configuration form appears on the same screen as the *BGP* menu.

Enable BGP

Enables BGP.

Autonomous System ID

Synopsis: *unsigned integer*

Autonomous System ID.

Always compare MED

Always comparing MED from different neighbors.

Default Local Preference

Synopsis: *unsigned integer*

Default: *100*

Default local preference value.

Deterministic Med

Pick the best-MED path among paths advertised from neighboring AS.

Router ID

Synopsis: *IPv4 address in dotted-decimal notation*

Router ID.

Figure 34.17. Distance Form

The path to the Distance form is *routing/dynamic/bgp*. Distance represents the distance value of BGP.

External Routes Distance

Synopsis: *unsigned integer*

Distance value for external routes.

Internal Routes Distance

Synopsis: *unsigned integer*

Distance value for internal routes.

Local Routes Distance

Synopsis: *unsigned integer*

Distance value for local routes.

34.5.1.1. Filter

The following information is available in configuration forms and display tables under the Filter submenu.

Autonomous System Path Filter

Name

Synopsis: *A string conforming to: "[^\s]+"*

Name of the AS-path filter.

Autonomous System Path Entry

Action

Synopsis: *string - one of the following keywords { permit, deny }*

Action.

Match

Synopsis: *string*

The regular expression to match the BGP AS paths.

Prefix List

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The name of the prefix list.

Description

Synopsis: *string*

The description of the prefix list.

Prefix List Entry

Sequence Number

Synopsis: *unsigned integer*

Sequence number of the entry.

Action

Synopsis: *string - one of the following keywords { permit, deny }*

Default: *permit*

Action.

Network

Synopsis: *IPv4 address and prefix in CIDR notation*

Network (xxx.xxx.xxx.xxx/xx).

Less Than or Equal to

Synopsis: *unsigned byte integer*

The maximum prefix length to be matched.

Greater Than or Equal to

Synopsis: *unsigned byte integer*

The minimum prefix length to be matched.

Route Map

Route Map Tag

Synopsis: *A string conforming to: "[^\s]+"*

Route map tag.

Route Map Entry

Sequence Number

Synopsis: *unsigned integer*

The sequence number of the route-map entry.

Action

Synopsis: *string - one of the following keywords { permit, deny }*

Default: *permit*

Action.

Call Route Map

Synopsis: *A string conforming to: "[^\s]+"*

Jump to another route-map after match+set.

On Match Goto

Synopsis: *unsigned integer*

Go to this entry on match.

Route Map Match

AS Path Filter

Synopsis: *A string conforming to: "[^\s]+"*

Match the BGP AS path filter.

Match Address of Route

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The prefix list name.

Match Next-Hop of Route

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The prefix list name.

Route Source Match

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The prefix list name.

Route Map Metric

Metric

Synopsis: *unsigned integer*

Match the route metric.

Peer Address

Synopsis: *IPv4 address in dotted-decimal notation*

Match the peer address.

Origin

Synopsis: *string - one of the following keywords { incomplete, igp, egp }*

Match the BGP origin code.

Aggregator

AS Number

Synopsis: *unsigned integer*

AS number.

IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of aggregator.

Prepend

AS Number

Synopsis: *unsigned integer*

AS number.

Exclude

AS Number

Synopsis: *unsigned integer*

AS number.

Local Preference

Synopsis: *unsigned integer*

Local preference.

Metric

operation

Synopsis: *string - one of the following keywords { sub, add, set }*

Set , add or subtract the metric value.

value

Synopsis: *unsigned integer*

Value.

next-hop

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *string - the keyword { peer }*

The next hop address (xxx.xxx.xxx.xxx/xx or peer to use peer address).

origin

Synopsis: *string - one of the following keywords { incomplete, igp, egp }*

The origin code.

originator-id

Synopsis: *IPv4 address in dotted-decimal notation*

The originator ID.

weight

Synopsis: *unsigned integer*

Weight.

34.5.1.2. Network

Networks may be specified in order to add BGP routers connected to the specified subnets. Note that a network specification need not be a given valid entry in the routing table. Since BGP is a border gateway protocol, one would more typically enter a more general network specification here. For example, if a routed network inside the AS comprised many different Class C subnets (/24) of the 192.168.0.0/16 range, it would be more efficient to advertise the one Class B network specification, 192.168.0.0/16, to one's BGP neighbors.

If BGP Neighbors are specified but no Networks are specified, then the router will receive BGP routing information from its neighbors but will not advertise any routes to them.

Subnet Address/Prefix

Synopsis: *IPv4 address and prefix in CIDR notation*

IP address/prefix.

34.5.1.3. Neighbor

Neighbors are other BGP routers with which to exchange routing information. One or more neighbors must be specified in order for BGP to operate.

> *If BGP Neighbors are specified but no Networks are specified, then the router will receive BGP routing information from its neighbors but will not advertise any routes to them.*

Neighbor IP address

Synopsis: *IPv4 address in dotted-decimal notation*

The neighbor IP address.

Neighbor Autonomous System ID

Synopsis: *unsigned integer*

A BGP neighbor.

ebgp-multihop

Synopsis: *unsigned byte integer*

The maximum hop count. This allows EBGP neighbors not on directly connected networks.

Maximum Prefix

Synopsis: *unsigned integer*

The maximum prefix number accepted from this peer.

next-hop-self

Disables the next hop calculation for this neighbor.

password

Synopsis: *A string conforming to: "[^\s]+"*

Password.

soft-reconfiguration

Per neighbor soft reconfiguration.

weight

Synopsis: *unsigned short integer*

The default weight for routes from this neighbor.

Route Map

in

Synopsis: *A string conforming to: "[^\s]+"*

Apply route map to incoming routes

out

Synopsis: *A string conforming to: "[^\s]+"*

Apply route map to outbound routes.

34.5.1.4. Aggregate-address

Aggregate Network

subnet

Synopsis: *IPv4 address and prefix in CIDR notation*

subnet (xxx.xxx.xxx.xxx/xx).

Options

value

Synopsis: *string - one of the following keywords { summary-only, as-set }*

Aggregate address option.

34.5.1.5. Distance-ip

Distance with Matched Subnet

Subnet/Prefix

Synopsis: IPv4 address and prefix in CIDR notation

IP Address/Prefix.

Distance

Synopsis: unsigned integer

Distance value.

34.5.1.6. Redistribute

Redistribute Route from Other Protocols

Redistribute Route From

Synopsis: *string - one of the following keywords { rip, ospf, connected, static, kernel }*

Redistribute route type.

Metric

Synopsis: *unsigned integer*

Metric value for redistributed routes.

35. Static Routing

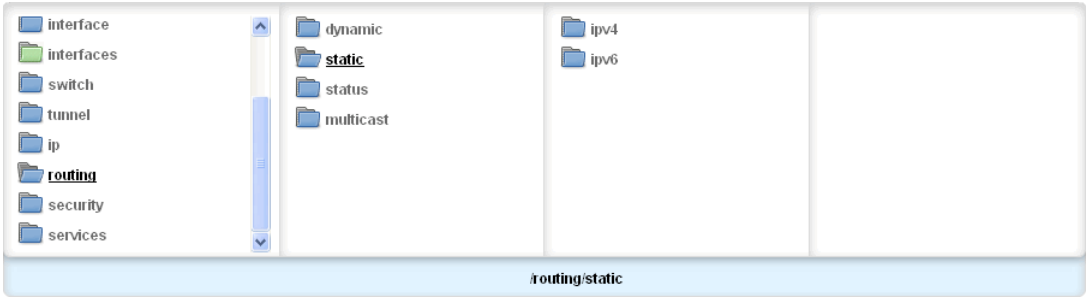

Figure 35.1. Static Menu

The Static menu is accessible from the main menu under *routing*. The path to this menu is *routing/static*.

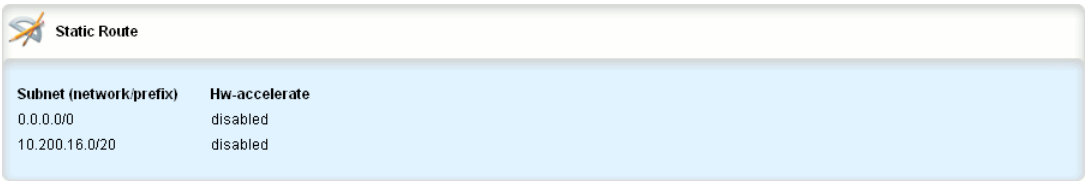

Figure 35.2. Static Route table

The path to the Static Route table is *routing/static/ipv4*.

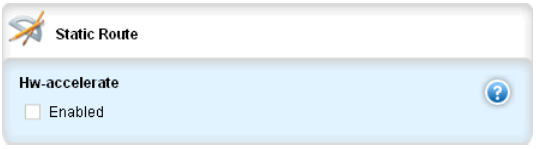

Figure 35.3. Static Route form

The path to the Static Route form is *routing/static/ipv4/{route}*.

hw-accelerate

If the static unicast route can be hardware accelerated, the option will be available. For a static unicast route to be accelerated, the ingress and egress interfaces must be switched.

The Hw-accelerate field will only be visible in the form if the device has a layer 3 switch.

Figure 35.4. Static Route Using Gateway table

The path to the Static Route Using Gateway table is *routing/static/ipv4/{route}/via*.

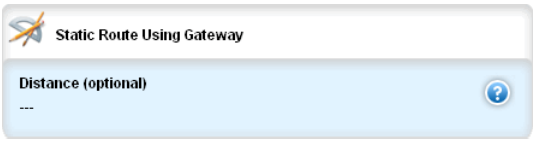

Figure 35.5. Static Route Using Gateway form

The path to the Static Route Using Gateway form is *routing/static/ipv4/{route}/via/{address}*.

Gateway Address

Synopsis: *IPv4 address in dotted-decimal notation*

The gateway for the static route.

Distance (optional)

Synopsis: *unsigned integer*

The distance for the static route.

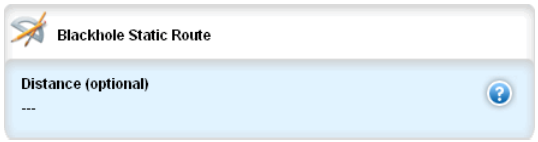

Figure 35.6. Blackhole Static Route form

The path to the Blackhole Static Route form is *routing/static/ipv4/{route}/blackhole*.

Distance (optional)

Synopsis: *unsigned integer*

The distance for the static route.

Figure 35.7. Static Route Using Interface table

The path to the Static Route Using Interface table is *routing/static/ipv4/{route}/dev*.

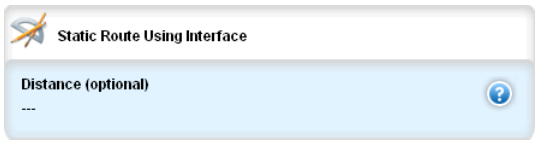

Figure 35.8. Static Route Using Interface form

The path to the Static Route Using Interface form is *routing/static/ipv4/{route}/dev/{interface}*.

Interface Name

Synopsis: *A string*

The interface for the static route.

Distance (optional)

Synopsis: *unsigned integer*

The distance for the static route.

36. Routing Status

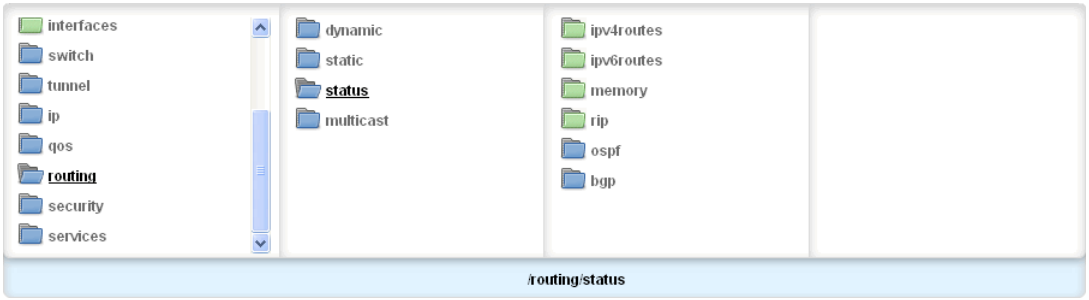

Figure 36.1. Routing Status Menu

The Routing Status menu is accessible under *routing/status*.

36.1. IPv4

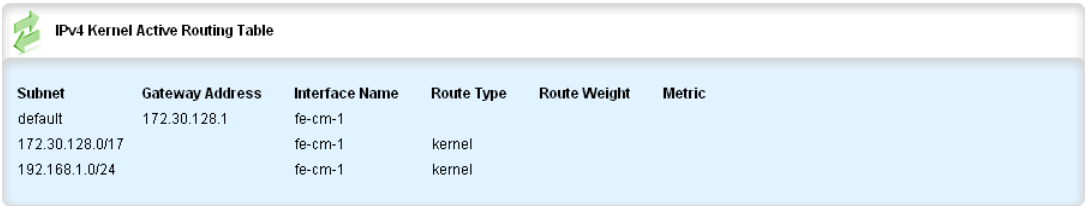

Figure 36.2. IPv4 Kernel Active Routing Table

The path to the IPv4 Kernel Active Routing table is *routing/status/ipv4routes*.

> *It is possible to create a route on a locally connected broadcast network (i.e. without a gateway) without also bringing up a corresponding IP address on that interface. For example, it would be possible to add* ***`192.168.1.0/24`*** *to* ***`switch.0001`****, which has an IP address of* ***`10.0.1.1`*** *but no corresponding alias address on the* ***`192.168.1.0/24`*** *subnet.*

Subnet

Synopsis: *string*

The network/prefix.

Gateway Address

Synopsis: *string*

The gateway address.

Interface Name

Synopsis: *string*

The interface name.

Route Type

Synopsis: *string*

The route type.

Route Weight

Synopsis: *string*

The route weight.

Metric

Synopsis: *string*

The route metric value.

36.2. IPv6

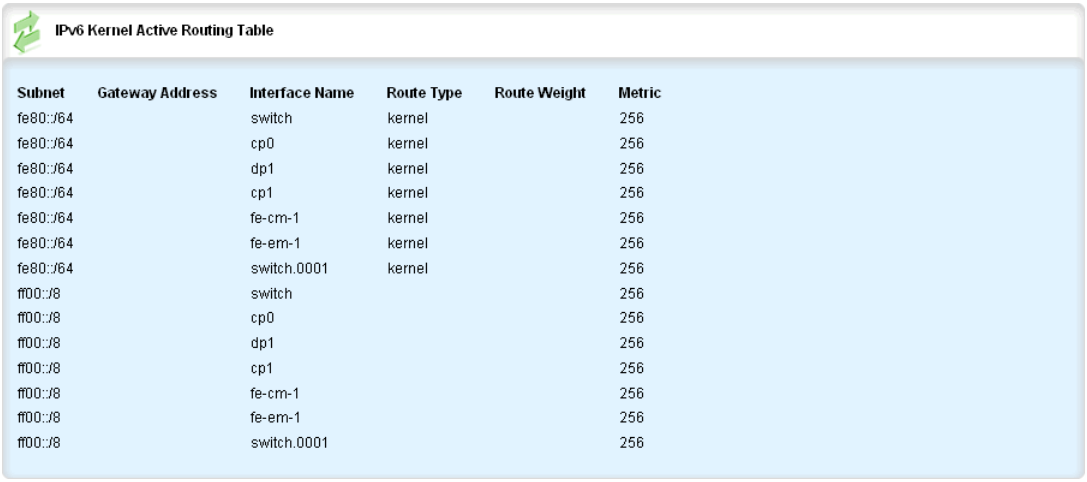

IPv6 Kernel Active Routing Table

Subnet	Gateway Address	Interface Name	Route Type	Route Weight	Metric
fe80::/64		switch	kernel		256
fe80::/64		cp0	kernel		256
fe80::/64		dp1	kernel		256
fe80::/64		cp1	kernel		256
fe80::/64		fe-cm-1	kernel		256
fe80::/64		fe-em-1	kernel		256
fe80::/64		switch.0001	kernel		256
ff00::/8		switch			256
ff00::/8		cp0			256
ff00::/8		dp1			256
ff00::/8		cp1			256
ff00::/8		fe-cm-1			256
ff00::/8		fe-em-1			256
ff00::/8		switch.0001			256

Figure 36.3. IPv6Kernel Active Routing Table

The path to the IPv6 Kernel Active Routing table is *routing/status/ipv6routes*.

Subnet

Synopsis: *string*

The network/prefix.

Gateway Address

Synopsis: *string*

The gateway address.

Interface Name

Synopsis: *string*

The interface name.

Route Type

Synopsis: *string*

The route type.

Route Weight

Synopsis: *string*

The route weight.

Metric

Synopsis: *string*

The metric value.

36.3. Memory Statistics

The path to the memory forms is *routing/status/memory*.

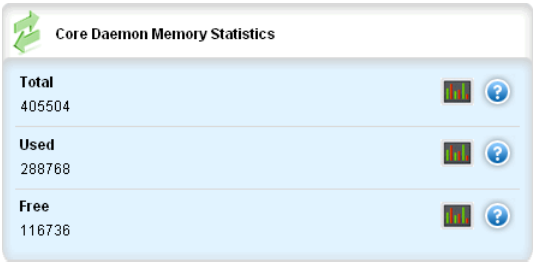

Figure 36.4. Core Daemon Memory Statistics Form

Total heap allocated (Byte)

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

Used ordinary blocks (Byte)

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

Free ordinary blocks (Byte)

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

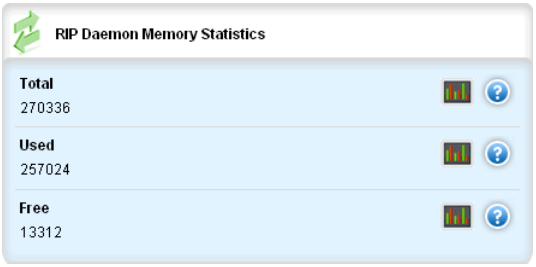

Figure 36.5. RIP Daemon Memory Statistics Form

total

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

used

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

free

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

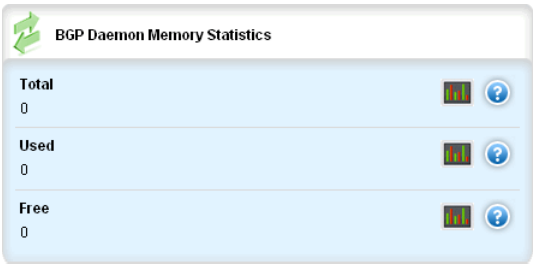

Figure 36.6. BGP Daemon Memory Statistics Form

total

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

used

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

free

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

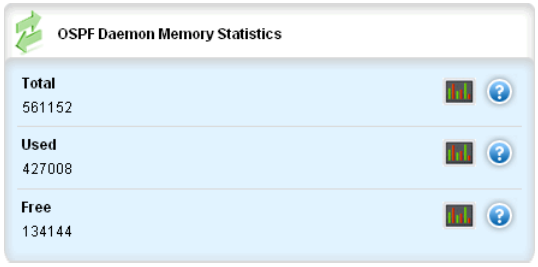

Figure 36.7. OSPF Daemon Memory Statistics Form

total

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

used

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

free

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

36.4. RIP

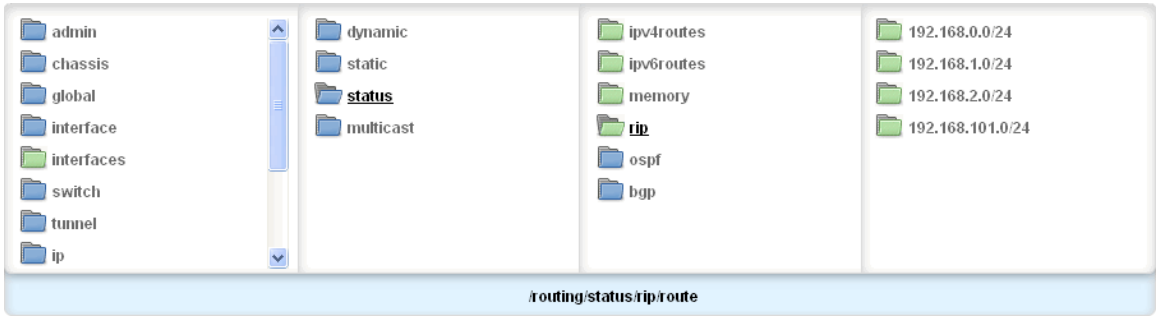

Figure 36.8. RIP Menu

The RIP status tables and forms are accessible from the RIP menu at *routing/status/rip* and *routing/status/rip/{address}*.

36.5. OSPF

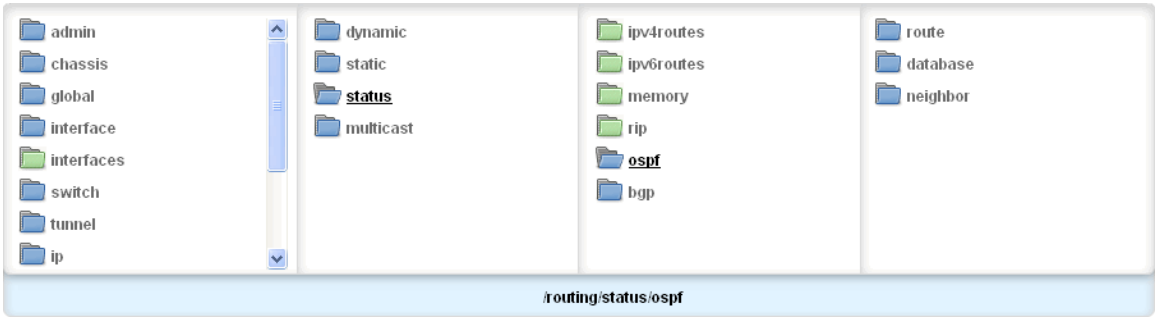

Figure 36.9. OSPF Menu

To display the OSPF menu, navigate to *routing/status/ospf*.

/routing/status/ospf/route/network

Id	Discard	Inter-area	Cost	Area
2.0.0.0/24	no	intra area	12	0.0.0.0
2.0.1.0/24	no	inter area	12	0.0.0.0
192.168.2.0/24	no	intra area	10	0.0.0.0
192.168.22.0/24	no	intra area	10	0.0.0.1

Figure 36.10. Network Table

To display the Network table, navigate to *routing/status/ospf/route/network*.

id

Synopsis: *string*

Network Prefix.

discard

Synopsis: *string*

This entry is discarded entry.

inter-area

Synopsis: *string*

Is path type inter area.

cost

Synopsis: *string*

Cost.

area

Synopsis: *string*

Area.

/routing/status/ospf/route/network/reach

How
via 192.168.2.1, switch.0021

Figure 36.11. Reach Table

To display the Reach table, navigate to *routing/status/ospf/route/network/{address}/reach*.

how

Synopsis: *string*

How to reach this network.

/routing/status/ospf/route/router

Id
1.1.1.1
2.0.0.1

Figure 36.12. Router Table

To display the Router table, navigate to *routing/status/ospf/route/router*.

id

Synopsis: *string*

Router ID.

/routing/status/ospf/route/router/area

Id	Inter-area	Cost	Abr	Asbr
0.0.0.1	intra area	10	no	yes

Figure 36.13. Area Table

To display the Area table, navigate to *routing/status/ospf/route/router/{number}/area*.

id

Synopsis: *string*

Area ID.

inter-area

Synopsis: *string*

Is path type inter area.

cost

Synopsis: *string*

Cost.

abr

Synopsis: *string*

Is area border router.

asbr

Synopsis: *string*

Is autonomous system border router.

/routing/status/ospf/database/net

Id	Area	Adv-router	Age	Seqnum
192.168.0.1	0.0.0.0	192.168.2.1	2632	0x8000009b

Figure 36.14. Net Table

To display the Net table, navigate to *routing/status/ospf/database/net*.

id

Synopsis: *string*

Link ID.

area

Synopsis: *string*

Area ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

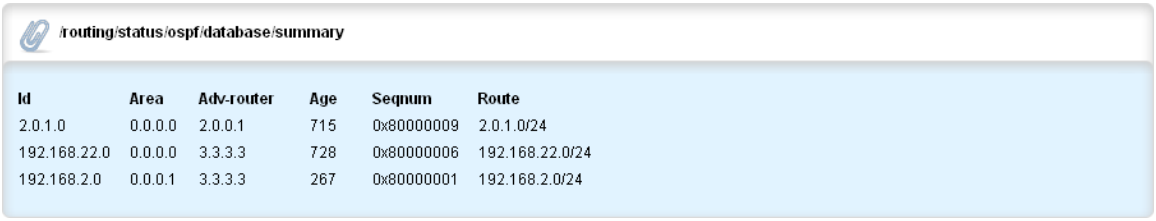

Figure 36.15. Summary Table

To display the Summary table, navigate to *routing/status/ospf/database/summary*.

id

Synopsis: *string*

Link ID.

area

Synopsis: *string*

Area ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

route

Synopsis: *string*

Route.

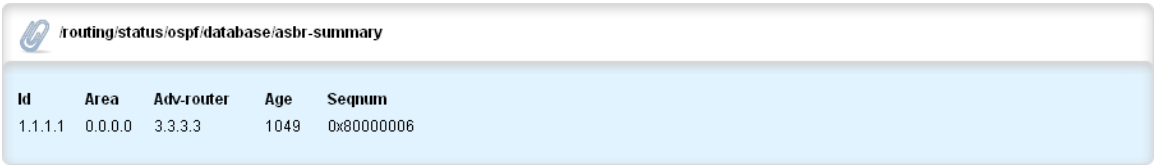

Figure 36.16. ASBR-Summary Table

To display the ASBR-Summary table, navigate to *routing/status/ospf/asbr-summary*.

id

Synopsis: *string*

Link ID.

area

Synopsis: *string*

Area ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

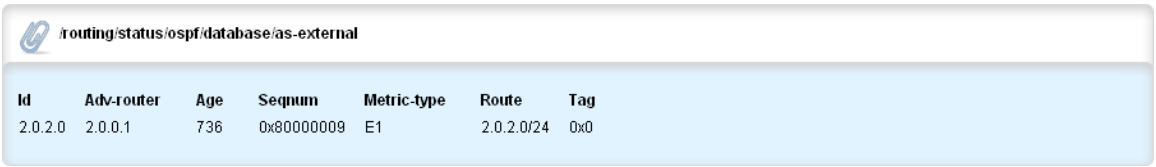
/routing/status/ospf/database/as-external
Id Adv-router Age Seqnum Metric-type Route Tag
2.0.2.0 2.0.0.1 736 0x80000009 E1 2.0.2.0/24 0x0

Figure 36.17. AS-External Table

To display the AS-External table, navigate to *routing/status/ospf/database/as-external*.

id

Synopsis: *string*

Link ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

metric-type

Synopsis: *string*

External metric type.

route

Synopsis: *string*

Route.

tag

Synopsis: *string*

Route tag.

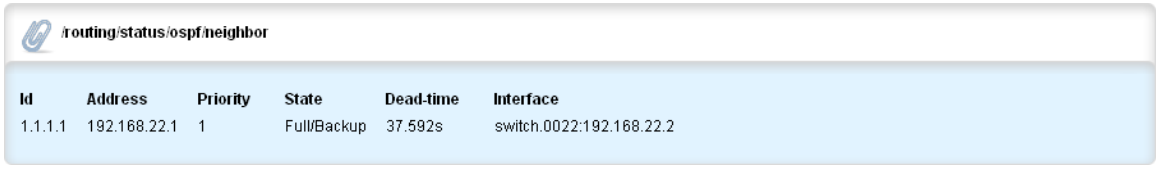

Figure 36.18. Neighbor Table

To display the Neighbor table, navigate to *routing/status/ospf/neighbor*.

id

Synopsis: *string*

Neighbor ID.

address

Synopsis: *string*

Address.

priority

Synopsis: *integer*

Priority.

state

Synopsis: *string*

State.

dead-time

Synopsis: *string*

Dead Time.

interface

Synopsis: *string*

Interface.

36.6. BGP

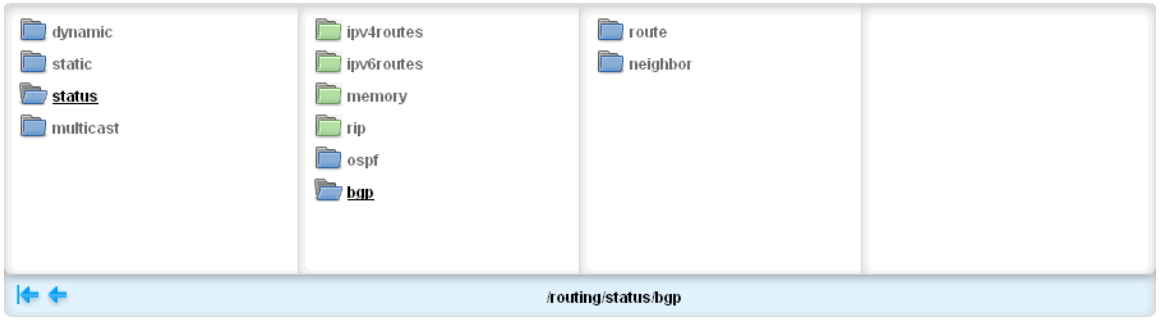

Figure 36.19. BGP Menu

To display the BGP menu, navigate to *routing/status/bgp*.

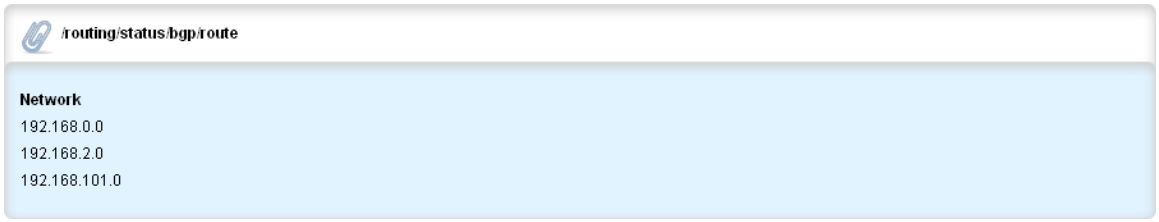

Figure 36.20. Route Table

To display the BGP Route table, navigate to *routing/status/bgp/route*.

network

Synopsis: *string*

Network.

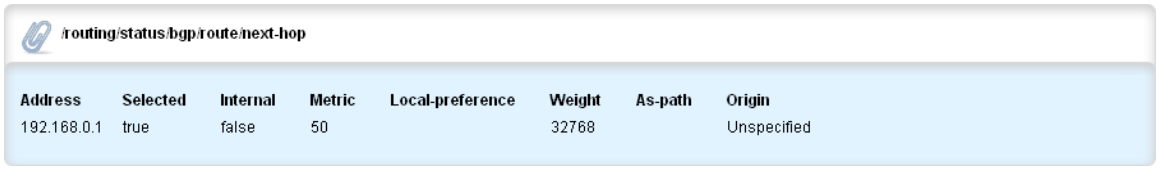
/routing/status/bgp/route/next-hop
Address Selected Internal Metric Local-preference Weight As-path Origin
192.168.0.1 true false 50 32768 Unspecified

Figure 36.21. Next Hop Table

To display the Next Hop table, navigate to *routing/status/bgp/route/{address}/next-hop*.

address

Synopsis: *string*

Next-hop address.

selected

Synopsis: *boolean*

Selected next-hop for this route.

internal

Synopsis: *boolean*

Internal route.

metric

Synopsis: *string*

Metric.

local-preference

Synopsis: *string*

Local preference.

weight

Synopsis: *integer*

Weight.

as-path

Synopsis: *string*

AS path.

origin

Synopsis: *string*

Origin.

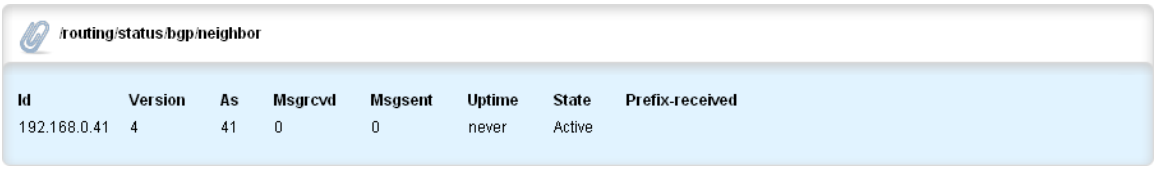

Figure 36.22. BGP Neighbor Table

To display the Neighbor table, navigate to *routing/status/bgp/neighbor*.

id

Synopsis: *string*

Neighbor address.

version

Synopsis: *integer*

BGP version.

as

Synopsis: *string*

Remote AS number.

msgrcvd

Synopsis: *integer*

Number of received BGP messages.

msgsent

Synopsis: *integer*

Number of sent BGP messages.

uptime

Synopsis: *string*

Peer up time.

state

Synopsis: *string*

Connection state with this neighbor.

prefix-received

Synopsis: *string*

Number of prefixes (networks) received from this neighbor.

37. Multicast Routing

Figure 37.1. Multicast Routing menu

The Multicast Routing menu is accessible from the main menu under *routing*. The path to this menu is *routing/multicast*. The user can choose between enabling dynamic multicast routing or static multicast routing by checking off "Enable" on the Routing Multicast Dynamic Form or the Routing Multicast Static Form.

Figure 37.2. Static Multicast Routing Configuration form

The Static Multicast Routing Configuration form appears on the same screen as the *Multicast* menu.

enabled

Enables static multicast routing service

Figure 37.3. Static menu

The path to the Static menu is *routing/multicast/static*. From the *static* menu, there are two branches of menus. By clicking on the *mcast-groups* menu, the user can access forms used to create multicast groups for forwarding. Clicking on the *status* menu allows you to view the status of your configured multicast groups. The following forms and tables are linked to the *mcast-groups* menu. The features available from the *status* menu are covered a little later in this chapter.

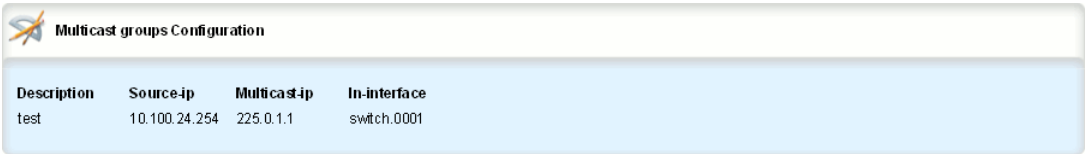

Figure 37.4. Multicast Groups Configuration table

The path to the Multicast Groups Configuration table is *routing/multicast/static/mcast-groups*.

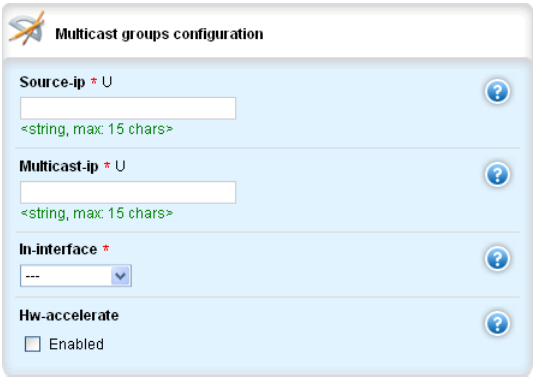

Figure 37.5. Multicast Groups Configuration form

The path to the Multicast Groups Configuration form is *routing/multicast/static/mcast-groups* and then clicking on one of the linked submenus.

description

Synopsis: *A string*

Describes this multicast group

source-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The expected source IP address of the multicast packet, in the format xxx.xxx.xxx.xxx

“U” indicates that this address is uniquely paired with the multicast address set in the **Multicast-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Multicast-ip address.

multicast-ip

Synopsis: *A string conforming to: "(22[4-9]|23[0-9])\.(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){2}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])"*

The multicast IP address to be forwarded, in the format xxx.xxx.xxx.xxx

The address must be in the range of 224.0.0.0 to 239.255.255.255

“U” indicates that this address is uniquely paired with the source IP address set in the **Source-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Source-ip address.

in-interface

Synopsis: *A string*

The interface upon which the multicast packet arrives

hw-accelerate

If the multicast route can be hardware accelerated, the option will be available. For a multicast route to be accelerated, the ingress and egress interfaces must be switched.

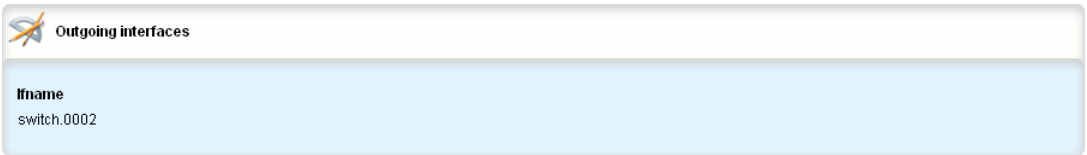

Figure 37.6. Outgoing Interfaces table

The path to the Outgoing Interfaces table is *routing/multicast/static/mcast-groups* and then clicking on one of the linked submenus.

The OutInterface is the interface to which the matched multicast packet will be forwarded.

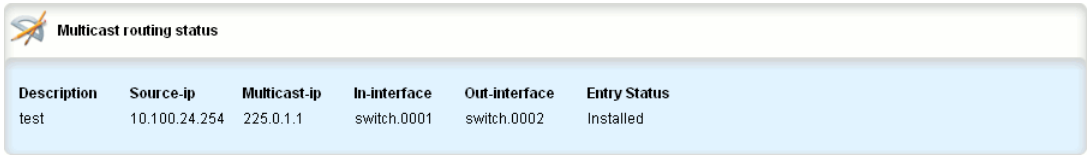

Figure 37.7. Multicast Routing Status table

The path to the Multicast Routing Status table is *routing/multicast/static/status*.

Figure 37.8. Multicast Routing Status form

The path to the Multicast Routing Status form is *routing/multicast/static/status* and then clicking on one of the linked submenus.

description

Synopsis: *A string*

Describes this multicast group

source-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The expected source IP address of the multicast packet, in the format xxx.xxx.xxx.xxx

“U” indicates that this address is uniquely paired with the multicast address set in the **Multicast-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Multicast-ip address.

multicast-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The multicast IP address to be forwarded, in the format xxx.xxx.xxx.xxx

The address must be in the range of 224.0.0.0 to 239.255.255.255

“U” indicates that this address is uniquely paired with the source IP address set in the **Source-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Source-ip address.

in-interface

Synopsis: *A string*

The interface upon which the multicast packet arrives

out-interface

Synopsis: *string*

The interface to which the matched multicast packet will be forwarded

entryStatus

Synopsis: *string*

The status of the multicast routing entry

38. Firewall

38.1. Firewall Fundamentals

Firewalls are software systems designed to prevent unauthorized access to or from private networks. Firewalls are most often used to prevent unauthorized Internet users from accessing private networks (intranets) connected to the Internet.

When the ROX™ firewall is used, the router serves as a *gateway* machine through which all messages entering or leaving the intranet pass. The router examines each message and blocks those that do not meet the specified security criteria. The router also acts as a *proxy*, preventing direct communication between computers on the Internet and intranet. Proxy servers can filter the kinds of communication that are allowed between two computers and perform address translation.

38.1.1. Stateless vs Stateful Firewalls

Firewalls fall into two broad categories: *stateless* and *stateful* (session-based).

Stateless or “static” firewalls make decisions about traffic without regard to traffic history; they simply open a "hole" for the traffic’s type, based upon a TCP or UDP port number. Stateless firewalling is relatively simple, easily handling web and email traffic. However, stateless firewalls have some disadvantages. All holes opened in the firewall are always open, and connections are not opened or closed based on outside criteria. Static IP filters offer no form of authentication.

Stateful firewalling adds considerable complexity to the firewalling process by tracking the state of each connection.

A stateful firewall also looks at and tests each packet, and the tests or “rules” may be modified depending on packets that have already been processed. This is called “connection tracking”. Stateful firewalls can also recognize that traffic on connected sets of TCP/UDP ports is from a particular protocol and manage it as a whole.

38.1.2. Linux® netfilter, iptables, and the Firewall

ROX™ employs a stateful firewall system known as *netfilter*, a subsystem of the Linux kernel that provides the ability to examine IP packets on a per-session basis.

The netfilter system uses rulesets, which are collections of packet classification rules that determine the outcome of the examination of a specific packet. The rules are defined by *iptables*, a generic table structure syntax and utility program for the configuration and control of netfilter.

ROX™ implements an IP firewall using a structured user interface to configure iptables rules and netfilter rulesets.

38.1.3. Network Address Translation

Network Address Translation (NAT) enables a LAN to use one set of IP addresses for internal traffic and a second set for external traffic. The netfilter NAT function makes all necessary IP address translations as traffic passes between the intranet and Internet. NAT is often referred to in Linux as IP Masquerading.

NAT itself provides a type of firewall by hiding internal IP addresses. More importantly, NAT enables a network to use more internal IP addresses. Since they are used internally only, there is no possibility of conflict with IP addresses used by other organizations. Typically, an internal network is configured to use one or more of the reserved address blocks described in RFC1918:

IP Network/Mask	Address Range
10.0.0.0/8	10.0.0.0 - 10.255.255.255
172.16.0.0/12	172.16.0.0 - 172.31.255.255

IP Network/Mask	Address Range
192.168.0.0/16	192.168.0.0 - 192.168.255.255

Table 38.1. RFC1918 Reserved IP Address Blocks

As a packet from a host on the internal network reaches the NAT gateway, its source address and source TCP/UDP port number are recorded. The address and port number is translated to the public IP address and an unused port number on the public interface. When the Internet host replies to the internal host's packet, it is addressed to the NAT gateway's external IP at the translation port number. The NAT gateway searches its tables and makes the opposite changes it made to the outgoing packet. NAT then forwards the reply packet to the internal host.

Translation of ICMP packets happens in a similar fashion, but without the source port modification.

NAT can be used in static and dynamic modes. Static NAT masks the private IP addresses by translating each internal address to a unique external address. Dynamic NAT translates all internal addresses to one or more external addresses.

38.1.4. Port Forwarding

Port forwarding, also known as redirection, allows traffic coming from the Internet to be sent to a host behind the NAT gateway.

Previous examples have described the NAT process when connections are made from the intranet to the Internet. In those examples, addresses and ports were unambiguous.

When connections are attempted from the Internet to the intranet, the NAT gateway will have multiple hosts on the intranet that could accept the connection. It needs additional information to identify the specific host to accept the connection.

Suppose that two hosts, 192.168.1.10 and 192.168.1.20 are located behind a NAT gateway having a public interface of 213.18.101.62. When a connection request for http port 80 arrives at 213.18.101.62, the NAT gateway could forward the request to either of the hosts (or could accept it itself). Port forwarding configuration could be used to redirect the requests to port 80 to the first host.

Port forwarding can also remap port numbers. The second host may also need to answer http requests. As connections to port 80 are directed to the first host, another port number (such as 8080) can be dedicated to the second host. As requests arrive at the gateway for port 8080, the gateway remaps the port number to 80 and forwards the request to the second host.

Finally, port forwarding can take the source address into account. Another way to solve the above problem could be to dedicate two hosts 200.0.0.1 and 200.0.0.2 and have the NAT gateway forward requests on port 80 from 200.0.0.1 to 192.168.1.10 and from 200.0.0.2 to 192.168.1.20.

38.2. Firewall Quick Setup

For users familiar with the firewall the following will serves as a reminder of how to build the firewall. New users may wish to read Section 38.3, "Firewall Terminology And Concepts" before continuing.

1. Logically partition your network into zones. Will you establish a DMZ? Will all Ethernet interfaces need to forward traffic to the public network? Which interfaces are to be treated in a similar fashion?
2. Assign your interfaces to the zones. If using T1/E1, have you created your T1/E1 interfaces prior to building the firewall?
3. Set the default policies for traffic from zone to zone to be as restrictive as possible. Has the local zone been blocked from connecting to the DMZ or firewall? Does the DMZ or firewall need to accept connections? Which connections should be dropped and which reset? What logs are kept?
4. How is the network interface IP assigned, i.e. dynamically or statically? Do hosts at the central site need to know the local address?

5. If your network interface IP is dynamically assigned, configure masquerading.
6. If your network interface IP is statically assigned, configure Source Network address Translation (SNAT). If a sufficient number of IP addresses are provided by the ISP, static NAT can be employed instead.
7. If your hosts must accept sessions from the Internet, configure the rules file to support Destination Network address Translation (DNAT). Which hosts need to accept connections, from whom and on which ports?
8. Configure the rules file to override the default policies. Have external connections been limited to approved IP address ranges. Have all but the required protocols been blocked?
9. If you are supporting a VPN, add additional rules.
10. Validate the configuration using the method outlined in Section 38.5.2, “Working with Firewall Configurations”.
11. Activate the firewall. It is recommended to run a port scan of the firewall after activation and verify that any defined logging is functioning as expected.

38.3. Firewall Terminology And Concepts

This section provides background on various firewall terms and concepts. References are made to the section where configuration applies.

38.3.1. Zones

A network zone is a collection of interfaces, for which forwarding decisions are made, for example:

Name	Description
net	The Internet
loc	Your Local Network
dmz	Demilitarized Zone
fw	The firewall itself
vpn1	IPSec connections on w1ppp
vpn2	IPSec connections on w2ppp

Table 38.2. Network Zones

New zones may be defined at any time. For example, if all of your Ethernet interfaces are part of the local network zone, disallowing traffic from the Internet zone to the local zone will disallow it to all Ethernet interfaces. If you wanted some interfaces (but not others) to access the Internet, you could create another zone.

38.3.2. Interfaces

ROX™ Firewall interfaces are simply the LAN and WAN interfaces available to the router. You must place each interface into a network zone.

If an interface supports more than one subnet, place the interface in zone ‘Any’ and use the zone hosts setup (see below) to define a zone for each subnet on the interface.

An example follows:

Interface	Zone
switch.0001	loc
switch.0002	loc
switch.0003	Any
switch.0004	dmz

Interface	Zone
w1ppp	net

Table 38.3. Interfaces

38.3.3. Hosts

ROX™ firewall hosts are used to assign zones to individual hosts or subnets, on an interface which handles multiple subnets. This allows the firewall to manage traffic being forwarded back out the interface it arrived on, but destined for another subnet. This is often useful for VPN setups to handle the VPN traffic separately from the other traffic on the interface which carries the VPN traffic. An example follows:

Zone	Interface	IP Address or Network
local	switch.0003	10.0.0.0/8
guests	switch.0003	192.168.0.0/24

Table 38.4. Hosts

38.3.4. Policy

Firewall policies are the default actions for connection establishment between different firewall zones. Each policy is of the form:

Source-zone Destination-zone Default-action

You can define a policy from each zone to each other. You may also use a wildcard zone of “all” to represent all zones.

The default action describes how to handle the connection request. There are six types of actions: ACCEPT, DROP, REJECT, QUEUE, CONTINUE and NONE. The first three are the most widely used and are described here.

When the *ACCEPT* policy is used, a connection is allowed. When the *DROP* policy is used, a request is simply ignored. No notification is made to the requesting client. When the *REJECT* policy is used, a request is rejected with an TCP RST or an ICMP destination-unreachable packet being returned to the client.

An example should illustrate the use of policies.

Source Zone	Destination Zone	Policy
loc	net	ACCEPT
net	all	DROP
all	all	REJECT

Table 38.5. Policies

The above policies will:

- Allow connection requests *only* from your local network to the Internet. If you want to allow requests from a ROX™ console to the Internet, add a policy of ACCEPT fw zone to the net zone.
- Drop (ignore) all connection requests from the Internet to your firewall or local network, and
- Reject all other connection requests.

Note that a client on the Internet probing the TCP/UDP ports will receive no responses and will not be able to detect the presence of the router. A host in the local network will fail to connect to the router, but will receive a notification.

Note that the order of policies is important. If the last rule of this example were entered first then no connections at all would be allowed.

38.3.5. Masquerading and SNAT

Masquerading and Source NAT (SNAT) are forms of dynamic NAT.

Masquerading substitutes a single IP address for an entire internal network. Use masquerading when your ISP assigns you an IP address dynamically at connection time.

SNAT substitutes a single address or range of addresses that have been assigned by your ISP. Use SNAT when your ISP assigns you one or more static IP addresses that you wish to use for one or more internal hosts.

Interface Subnet Address Protocol Port(s)

Interface is the outgoing (WAN or Ethernet) interface and is usually your Internet interface.

Subnet is the subnet that you wish to hide. It can be an interface name (such as switch.0001) or a subnetted IP address.

Address is an (optional IP) address that you wish to masquerade as.

The presence of the Address field determines whether masquerading or SNAT is being used. Masquerading is used when only Interface and Subnet are present. SNAT is used when Interface, Subnet and Address are present.

Protocol (optionally) takes on the name of protocols (e.g. tcp, udp) that you wish to masquerade.

Ports (optionally) takes on the ports to masquerade when protocol is set to tcp or udp. These can be raw port numbers or names as found in file /etc/services.

Some examples should illustrate the use of masquerading:

Rule	Interface	Subnet	Address	Protocol	Ports
1	switch.0001	switch.0002			
2	ppp+	switch.0002	66.11.180.161		
3	ppp+	192.168.0.0/24	66.11.180.161		
4	w1ppp	switch.0001	100.1.101.16		
5	w1ppp	switch.0001	100.1.101.16	tcp	smtp

Table 38.6. Masquerading Examples

1. In this masquerading rule, port switch.0002 is connected to the local network and switch.0001 is connected to a DSL modem. Traffic from the subnet handled by switch.0002 should be translated to whatever IP is assigned to the modem. Internet clients will not be able to determine the router's public address unless some form of dynamic DNS is employed.
2. In this SNAT rule, a static address of 66.11.180.161 is acquired from the ISP. Traffic from the subnet handled by switch.0002 should be translated to 66.11.180.161 as it sent to the Internet over ppp. The + at the end of "ppp+" causes the ROX™ firewall to match any ppp interface.
3. This example is much the same as the previous one only the subnet is explicitly described, and could include traffic from any of the Ethernet ports.
4. In this SNAT rule, traffic from the subnet handled by only port switch.0001 should be translated to 100.1.101.16 as it sent to the Internet on t1/e1 port w1ppp.
5. This example is much the same as the previous one excepting that only smtp from switch.0001 will be allowed.

38.3.6. Rules

The default policies can completely configure traffic based upon zones. But the default policies cannot take into account criteria such as the type of protocol, IP source/destination addresses and the need to perform special actions such as port forwarding. The firewall rules can accomplish this.

The ROX™ firewall rules provide exceptions to the default policies. In actuality, when a connection request arrives, the rules file is inspected first. If no match is found then the default policy is applied. Rules are of the form:

Action Source-Zone Destination-Zone Protocol Destination-Port Source-Port Original-Destination-IP Rate-Limit User-Group

Actions are ACCEPT, DROP, REJECT, DNAT, DNAT-, REDIRECT, REDIRECT-, CONTINUE, LOG and QUEUE. The DNAT-, REDIRECT-, CONTINUE, LOG and QUEUE actions are not widely used used and are not described here.

Action	Description
ACCEPT	Allow the connection request to proceed.
DROP	The connection request is simply ignored. No notification is made to the requesting client.
REJECT	The connection request is rejected with an RST (TCP) or an ICMP destination-unreachable packet being returned to the client.
DNAT	Forward the request to another system (and optionally another port).
REDIRECT	Redirect the request to a local tcp port number on the local firewall. This is most often used to “remap” port numbers for services on the firewall itself.

Table 38.7.

The remaining fields of a rule are as described below:

Rule Field	Description
Action	The action as described in the previous table.
Source-Zone	The zone the connection originated from.
Destination-Zone	The zone the connection is destined for.
Protocol	The tcp or udp protocol type.
Destination-Port	The tcp/udp port the connection is destined for.
Source-Port	The tcp/udp port the connection originated from.
Original-Destination-IP	The destination IP address in the connection request as it was received by the firewall.
Rate-Limit	A specification which allows the rate at which connections are made to be limited.

Table 38.8.

Some examples will illustrate the power of the rules file:

Rule	Action	Source-Zone	Destination-Zone	Protocol	Dest-Port	Source-Port	Original-Destination-IP
1	ACCEPT	net:204.18.45.0/24	fw				
2	DNAT	net	loc:192.168.1.3	tcp	ssh, http		
3	DNAT	net:204.18.45.0/24	loc:192.168.1.3	tcp	http	-	130.252.100.69
4	ACCEPT	fw	net	icmp			
5	ACCEPT	net:204.18.45.0/24	fw	icmp	8		

Table 38.9.

1. This rule accepts traffic to the firewall itself from the 204.18.45.0/24 subnet. If the default policy is to drop all requests from net to the firewall, this rule will only accept traffic from the authorized subnet.
2. This rule forwards all ssh and http connection requests from the Internet to local system 192.168.1.3.

3. This rule forwards http traffic from 204.18.45.0/24 (which was originally directed to the firewall at 130.252.100.69) to the host at 192.168.1.3 in the local zone. If the firewall supports another public IP address (e.g. 130.252.100.70), a similar rule could map requests to another host.
4. and 5. These rules allow the firewall to issue icmp requests to the Internet and to respond to icmp echo requests from the authorized subnet.

Each of the Source and Destination zones may use one of the defined zone names, or one may select "Other..." and specify a zone name in the text field to the right. Both Source and Destination may be further qualified using the Only hosts in zone with addresses fields. Multiple comma-separated subnet, IP, or MAC addresses may be specified in the following way:

- IP subnet: 192.168.1.0/24
- IP address: 192.168.1.1
- IP address range: 192.168.1.1-192.168.1.25
- MAC address: ~00-A0-C9-15-39-78

38.4. Configuring The Firewall And VPN

38.4.1. Policy-based Virtual Private Networking

Begin configuration by creating local, network and vpn zones. Identify the network interface that carries the encrypted IPSec traffic and make this interface part of zone “ANY” in the interfaces menu as it will be carrying both traffic for both zones.

Visit the Host menu and, for the network interface that carries the encrypted IPSec traffic, create a zone host with zone VPN, the correct subnet and the IPSec zone option checked. If you plan to have VPN tunnels to multiple remote sites ensure that a zone host entry exists for each (or collapse them into a single subnet). Create another zone host for the same interface with a network zone, using a wider subnet mask such as 0.0.0.0/0. It is important that the vpn zone be declared before the net zone since the more specific vpn zone subnet must be inspected first.

Host Zone	Interface	Subnet	IPSec Zone?
vpn	w1ppp	192.168.1.0/24	Yes
net	w1ppp	0.0.0.0/0	No

Table 38.10.

The IPSec protocol operates on UDP port 500 and using protocols Authentication Header (AH) and Encapsulating Security Payload (ESP) protocols. The firewall must accept this traffic in order to allow IPSec.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	fw	ah	
ACCEPT	net	fw	esp	
ACCEPT	net	fw	udp	500

Table 38.11.

IPSec traffic arriving at the firewall is directed to openswan, the IPSec daemon. Openswan then decrypts the traffic and forwards it back to the ROX™ firewall on the same interface that originally received it. You will also need a rule to allow traffic to enter from this interface.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	vpn	loc		

Table 38.12.

38.4.2. Virtual Private Networking to a DMZ

If the firewall is to pass the VPN traffic through to another device (e.g. a VPN device in a DMZ) then establish a DMZ zone and install the following rules.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	dmz	ah	
ACCEPT	net	dmz	esp	
ACCEPT	net	dmz	udp	500
ACCEPT	dmz	net	ah	
ACCEPT	dmz	net	esp	
ACCEPT	dmz	net	udp	500

Table 38.13.

38.5. Firewall Configuration

All firewall fields accept only alphanumeric characters, excluding spaces. Do not use punctuation or other special characters in these fields.

Figure 38.1. Security Menu

The Security menu is a top-level menu that is accessible from the main menu. Items used to configure network security can be accessed from this menu.

Figure 38.2. Firewall Description table

Name

Synopsis: *string*

Description

Synopsis: *string*

An optional description string

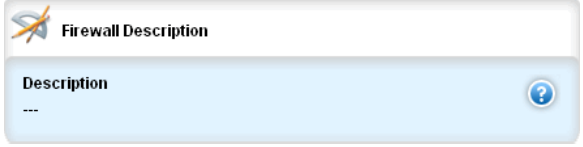

Figure 38.3. Firewall Description form

To display the Firewall form, navigate to *security/firewall* and then click on the submenu representing the configured firewall (for example, *firewall1*).

38.5.1. Adding a Firewall

To add a firewall, enter edit private mode, navigate to */security/firewall/fwconfig*, and click <Add fwconfig>.

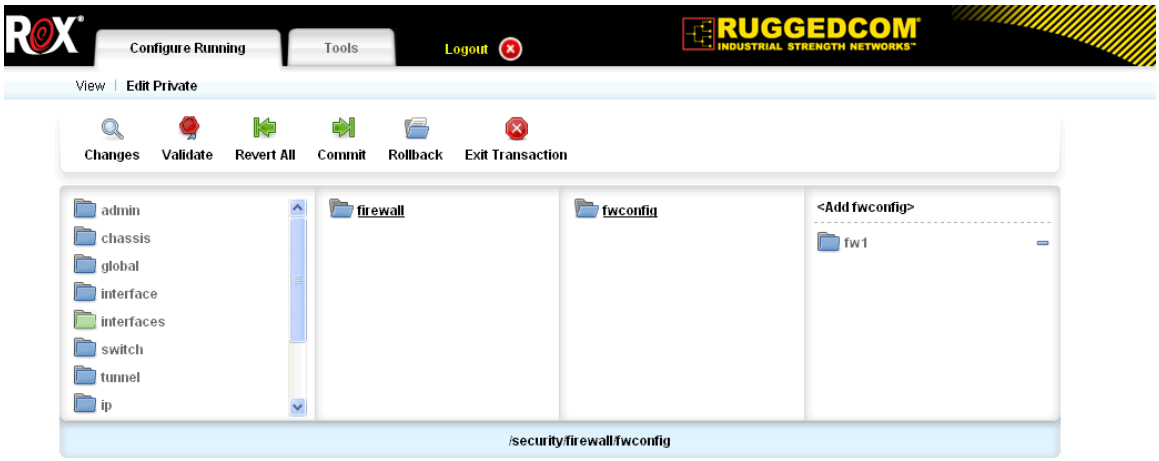

Figure 38.4. Adding a Firewall

In the *Key settings* form, enter a name for the firewall and click *Add* field.

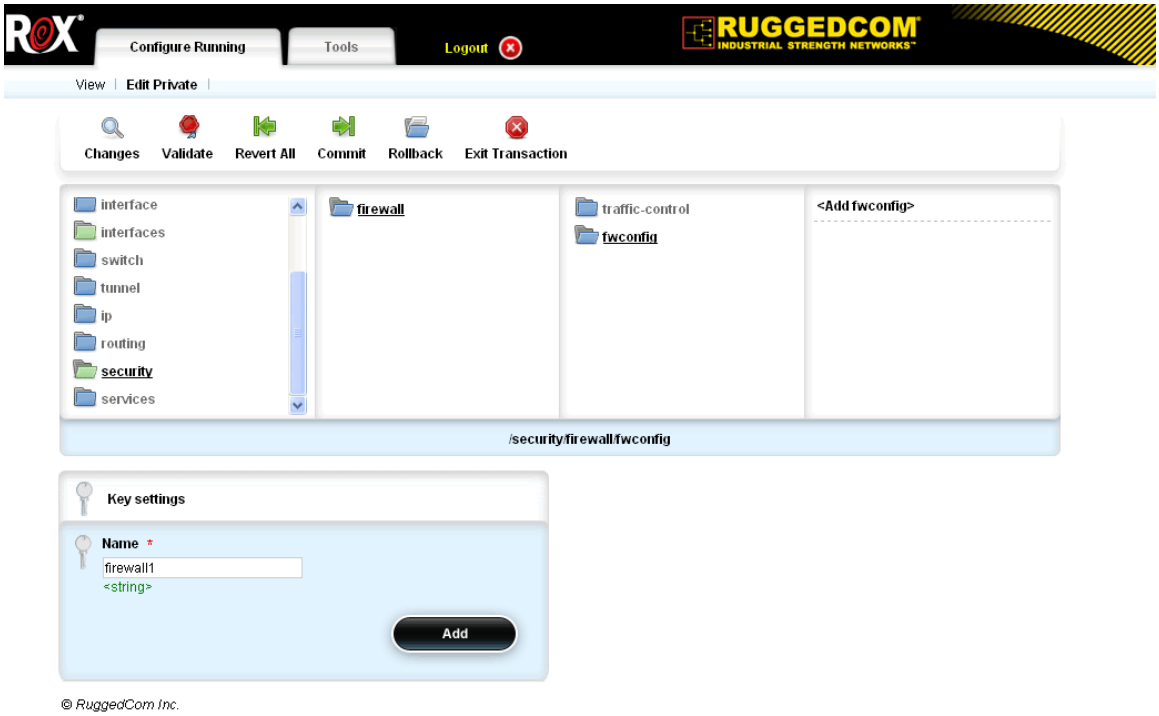

Figure 38.5. Naming a Firewall

After adding the firewall, the fwmasq, fwnat, fwrule, fwpolicy, fwinterface, fwhost, and fwzone submenus appear. To configure these areas, click on a submenu.

Figure 38.6. Firewall Submenus

38.5.2. Working with Firewall Configurations

The ROX™ firewall configuration system allows a network security administrator to work on one or more inactive firewall configurations while another is active and installed on the system. Section 38.5.2.1, "Typical Use Case" illustrates how to use the ROX™ firewall configuration system.

Control of the firewall configuration is achieved by using the three variables in the *Firewall Configuration* form, below:

Figure 38.7. Firewall Configuration form

Enable active configuration

Enables/disables the firewall configuration specified in active-config

Specify work configuration

Synopsis: *string*

The current work firewall is specified here.

Specify active configuration

Synopsis: *string*

The current active firewall is specified here

38.5.2.1. Typical Use Case

The following set of steps illustrates the configuration and maintenance of a set of firewall rules on an active ROX™ firewall system:

1. On an unconfigured system, begin configuring a set of firewall rules by giving the firewall a name: 'fw1', adding zones, interfaces, etc. At each commit at this stage, configuration data is saved but no validation is performed.
2. In order to validate the 'fw1' firewall configuration in progress, set the *work-config* variable to the name: 'fw1' and commit the changes. The system validates the firewall configuration named 'fw1' and displays the results. Note that the configuration in progress is saved whether or not the

validation succeeds. A configuration in progress may be validated in this way at any time without affecting an active firewall configuration.

3. After ‘fw1’ has been verified, it may be made active in the system by setting the *active-config* variable to the name: ‘fw1’, setting *firewall-enable* and committing the changes. The system validates the active firewall configuration and if it finds no errors, it activates the ‘fw1’ firewall configuration.
4. While the ‘fw1’ firewall configuration is active, you might wish to make changes without altering the live configuration. Using the CLI, copy the firewall configuration named ‘fw1’ to ‘fw2’. Change the *work-config* variable to ‘fw2’. Any configuration changes made to ‘fw2’ are validated when you commit your changes, and any errors in ‘fw2’ are displayed. An alternate configuration may be modified and validated in this way at any time without affecting an active firewall configuration.
5. Alternatively, while the ‘fw1’ firewall configuration is active, you might wish to make changes to the live configuration. Any changes made to a configuration that is defined as ‘active-config’ and ‘enable’ will be reflected on the live configuration currently running on the system, pending a successful validation. For instance, work-config can be a configuration named ‘fw2’ while active-config is ‘fw1’ and enabled. Modifying fw1 in this case will, upon successful validation, update the running configuration to reflect the changes.

38.5.3. Zone Configuration

Zones are collections of interfaces, for which forwarding decisions are made. They are made of different networks reachable from this system, defined by name and type of zone.

Figure 38.8. Zone table

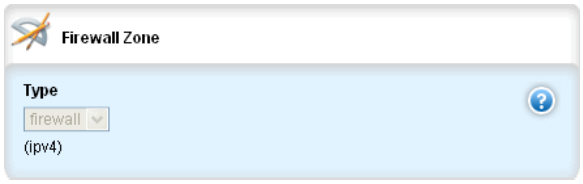

Figure 38.9. Zone form

This form allows you to add, delete and configure zones. New zones can be added by entering the Edit Private mode and then adding zones.

Name

Synopsis: *A string*

A unique name to assign to this zone. Be sure to *also create* a zone called **fw** of type *firewall*.

Type

Synopsis: *string - one of the following keywords { firewall, ipsec, ipv4 }*

Default: *ipv4*

Zone types are: firewall, ipv4 or ipsec

description

Synopsis: *string*

(Optional) The description string for this zone

38.5.4. Interface Configuration

Figure 38.10. Main Interface Settings table

interface

Synopsis: *string*

Currently active or not - add '+' for same interfaces: ppp+

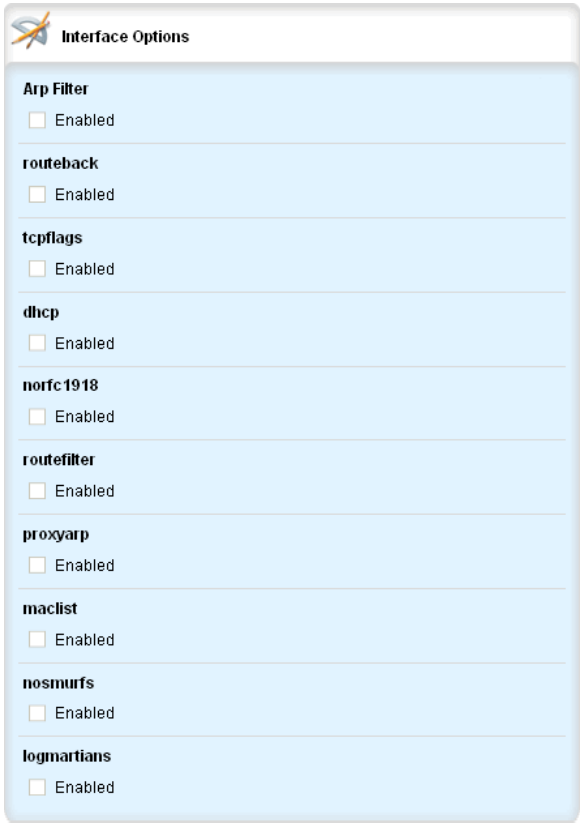

Figure 38.11. Interface Options form

Arp Filter

Responds only to ARP requests for configured IP addresses

routeback

Allow traffic on this interface to be routed back out that same interface

tcpflags

Illegal combinations of TCP flags dropped and logged at info level

dhcp

Allows DHCP datagrams to enter and leave the interface

norfc1918

Not currently implemented

routefilter

Enables route filtering

proxyarp

Enables proxy ARP

maclist

Not currently implemented

nosmurfs

Packets with broadcast address as source are dropped and logged at info level

logmartians

Enables logging of packets with impossible source addresses

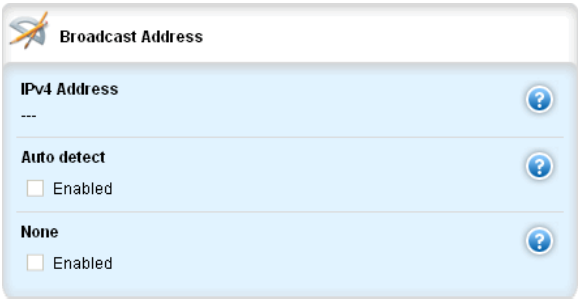

Figure 38.12. Broadcast Address form

broadcast-addr

(Optional) A broadcast address

38.5.5. Host Configuration

Hosts are used to assign zones to individual hosts or subnets, on an interface which handles multiple subnets.

Figure 38.13. Main Host Settings table

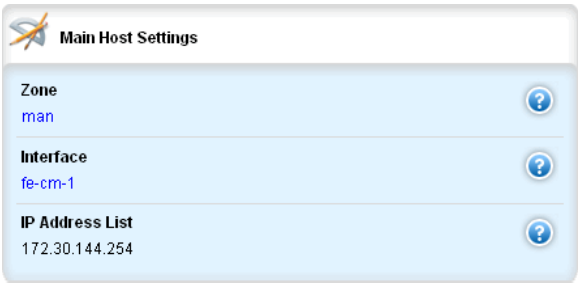

Figure 38.14. Main Host Settings form

name

Synopsis: *string*

The name of a host configuration entry

Zone

Synopsis: *A string*

A pre-defined zone

Interface

Synopsis: *A string*

A pre-defined interface to which optional IPs and/or networks can be added

IP Address List

Synopsis: *string*

(Optional) Additional IP addresses or networks - comma separated

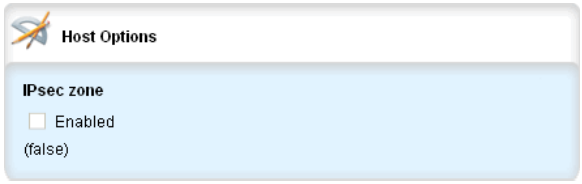

Figure 38.15. Host Options form

IPsec zone

Synopsis: *boolean*

Default: *false*

38.5.6. Policies

Figure 38.16. Main Policy Settings table

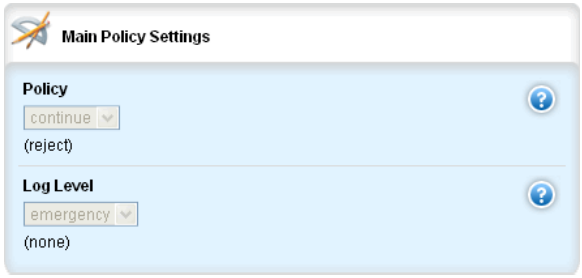

Figure 38.17. Main Policy Settings form

Default actions for connection establishment between different zones. Configuration of the default actions for traffic between different firewall zones. They can be overridden for particular hosts or types of traffic by defining specific rules.

Policy Name

Synopsis: *string*

Enter a name tag for this policy

Policy

Synopsis: *string - one of the following keywords { continue, reject, drop, accept }*

Default: *reject*

A default action for connection establishment between different zones.

Log Level

Synopsis: *string - one of the following keywords { emergency, alert, critical, error, warning, notice, info, debug, none }*

Default: *none*

(Optional) Whether or not logging will take place and at which logging level.

description

Synopsis: *string*

(Optional) The description string for this policy

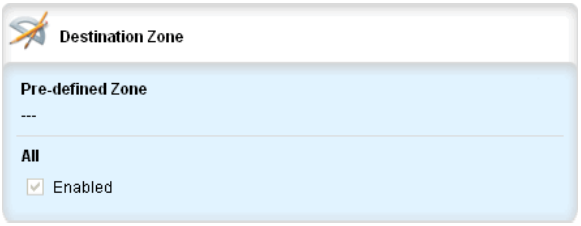

Figure 38.18. Destination Zone form

destination-zone

The zone in which the request terminates. Enter a destination zone configuration by specifiying a zone. Please choose either a pre-defined zone or all.

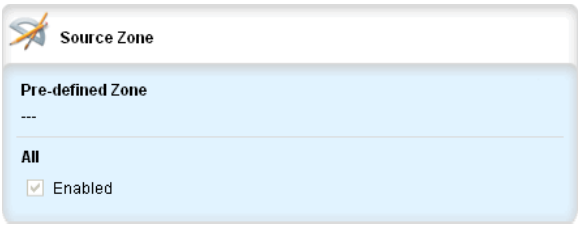

Figure 38.19. Source Zone form

source-zone

The zone from which the request originates. Enter a source zone configuration by specifiying a zone. Please choose either a pre-defined zone or all.

38.5.7. Network Address Translation

Configuration of one-to-one Network Address Translation (NAT). The static network address translation entries in this table can be used to set up a one-to-one correspondence between an external address on your firewall and an RFC1918 address of a host behind the firewall. Static NAT is often used to allow connections to an internal server from outside the network.

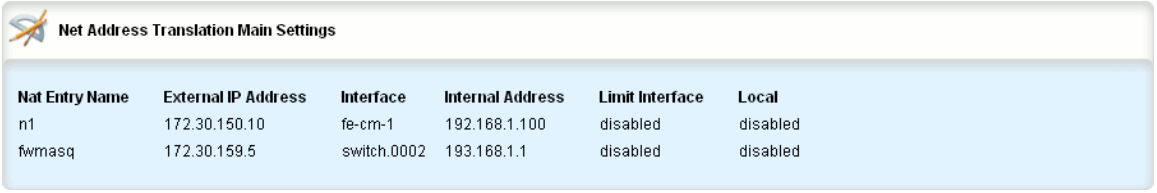

Figure 38.20. Net Address Translation Main Settings table

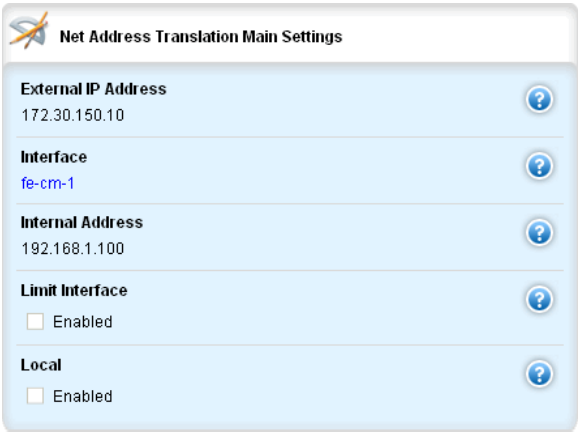

Figure 38.21. Net Address Translation Main Settings form

Nat Entry Name

Synopsis: *string*

Enter a name for this NAT entry

External IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

The external IP Address (must not be a DNS name)

Interface

Synopsis: *A string*

Interfaces that have the EXTERNAL address

Internal Address

Synopsis: *IPv4 address in dotted-decimal notation*

The internal address (must not be a DNS Name)

Limit Interface

Translation only effective from the defined interface (default: all interfaces)

Local

Translation effective from the firewall system (default: disabled)

description

Synopsis: *string*

(Optional) The description string for this nat entry

38.5.8. IP Masquerading

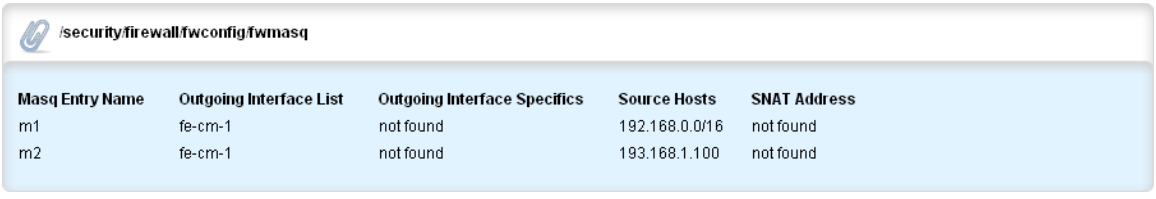

Figure 38.22. FWMasq table

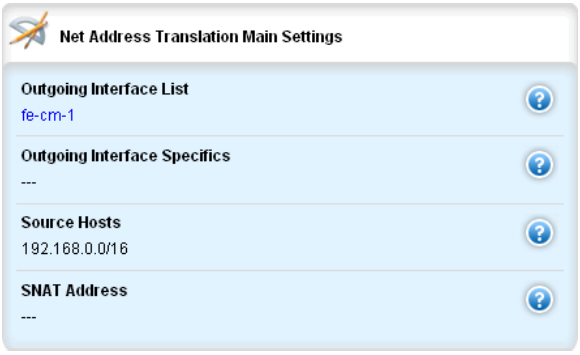

Figure 38.23. Net Address Translation Main Settings form

Masquerading substitutes a single IP address for an entire internal network

Masq Entry Name

Synopsis: *string*

A name for this masquerading configuration entry

Outgoing Interface List

Synopsis: *string*

An outgoing interfacelist - usually the internet interface

Outgoing Interface Specifics

Synopsis: *string*

(Optional) An outgoing interface list - specific destinations IP for the out-interface

Source Hosts

Synopsis: *string*

Subnet range or comma-separated list of hosts (IPs)

SNAT Address

Synopsis: *string*

(Optional) By specifying an address here, SNAT will be used and this will be the source address

description

Synopsis: *string*

(Optional) The description string for this masq entry

38.5.9. Rules

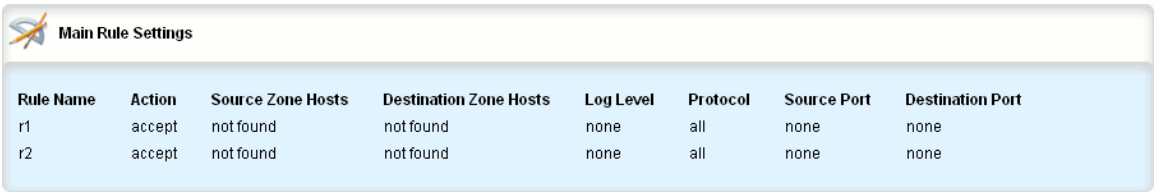
Main Rule Settings
Rule Name Action Source Zone Hosts Destination Zone Hosts Log Level Protocol Source Port Destination Port
r1 accept not found not found none all none none
r2 accept not found not found none all none none

Figure 38.24. Main Rule Settings table

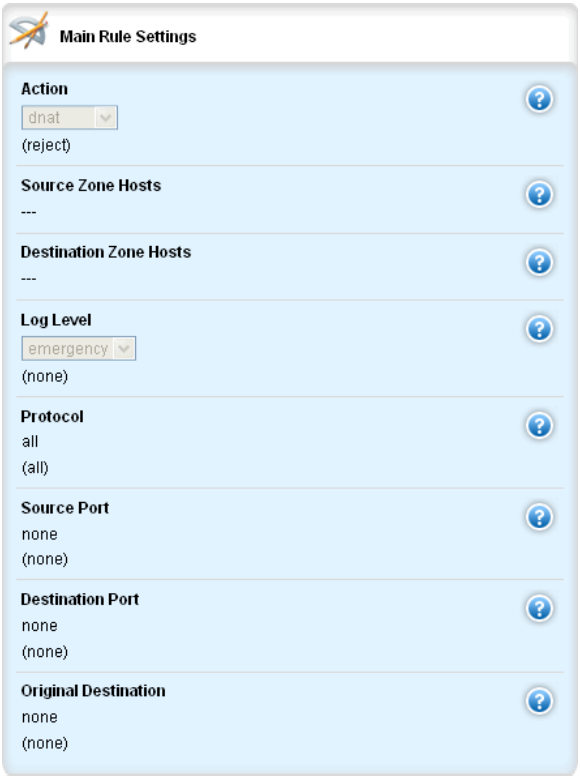

Figure 38.25. Main Rule Settings form

Rules are to establish exceptions to the default policies. This table lists exceptions to the default policies for certain types of traffic, sources or destinations. The chosen action will be applied to packets matching the chosen criteria instead of the default.

Rule Name

Synopsis: *string*

Enter a unique name that identifies this rule.

Action

Synopsis: *string - one of the following keywords { dnat, dnat-, redirect, continue, reject, drop, accept }*

Default: *reject*

The final action to take on incoming packets matching this rule.

Destination Zone Hosts

Synopsis: *string*

(Optional) Add comma-separated host IPs to the destination-zone - may include :port for DNAT or REDIRECT

Log Level

Synopsis: *string - one of the following keywords { emergency, alert, critical, error, warning, notice, info, debug, none }*

Default: *none*

(Optional) Whether or not logging will take place and at which logging level.

Protocol

Synopsis: *string*

Synopsis: *string - one of the following keywords { all, icmp, udp, tcp }*

Default: *all*

The protocol to match for this rule.

Source Port

Synopsis: *string*

Synopsis: *string - one of the following keywords { none, Related, Any }*

Default: *none*

(Optional) The tcp/udp port the connection originated from.

Destination Port

Synopsis: *string*

Synopsis: *string - one of the following keywords { none, Related, Any }*

Default: *none*

(Optional) The tcp/udp port the connection is destined for.

Original Destination

Synopsis: *string*

Synopsis: *string - the keyword { None }*

Default: *none*

(Optional) The destination IP address in the connection request as it was received by the firewall.

description

Synopsis: *string*

(Optional) The description string for this rule

Figure 38.26. Source Zone form

Source Zone Hosts

Synopsis: *string*

(Optional) Add comma-separated host IPs to a predefined source-zone

Figure 38.27. Destination Zone form

destination-zone

synopsis: string

(Optional) Add comma-separated host IPs to the destination-zone - may include :port for DNAT or REDIRECT

39. Traffic Control

Traffic Control (TC) is a firewall subsystem managing the amount of bandwidth per network interface that different types of traffic are permitted to use. For a traffic control configuration to work, a firewall must be configured.

A ROX™ system allows up to 4 different firewall configurations, enabling you to quickly change between configurations. You can quickly assess different configurations without needing to save and reload any part of the configuration. In contrast, there is only one traffic control configuration. When enabled, a traffic control configuration is used with the current firewall configuration. A current firewall configuration is defined as one that is specified in either work-config and/or active-config. It does not have to be enabled to be validated.

A TC configuration can be seen as an additional configuration that goes along with a current firewall configuration. To actually add a TC configuration, it has to be enabled by setting the /qos/traffic-control/ Basic or Advanced Configuration Modes variable. Only at that point will a TC configuration be added to the current firewall configuration.

On the RX1500 and RX5000, traffic control is not available on the Ethernet traffic on any line module when when Layer 3 hardware acceleration is enabled. Therefore, it is intended to be used only on the WAN interfaces.

39.1. Traffic Control Modes

Traffic Control functions are divided into two modes: *basic mode* and *advanced mode*. The basic mode contains functions with basic traffic control configuration parameters. The advanced mode contains functions with advanced traffic control configuration parameters. The two modes cannot be accessed simultaneously. Only the mode that is currently configured can be accessed.

39.1.1. Traffic Control Basic (basic-configuration) Configuration Mode

Basic-configuration mode offers a limited set of options and parameters. Use this mode to set the outgoing bandwidth for an interface, the interface priority (high, medium, low), and some simple traffic control characteristics. Basic traffic shaping affects traffic identified by protocol, port number, address, and interface. Note that some of these options are mutually exclusive; refer to the information given for each option.

In basic-configuration mode, a packet is categorized based on the contents of its TOS field if it does not match any of the defined bands.

39.1.2. Traffic Control Advanced (advanced-configuration) Configuration Mode

In advanced-configuration mode, each interface to be managed is assigned a total bandwidth that it should allow for incoming and outgoing traffic. Classes are then defined for each interface, each with its own minimum assured bandwidth and a maximum permitted bandwidth. The combined minimum of the classes on an interface must be no more than the total outbound bandwidth specified for the interface. Each class is also assigned a priority, and any bandwidth left over after each class has received its minimum allocation (if needed) will be allocated to the lowest priority class up until it reaches its maximum bandwidth, after which the next priority is allocated more bandwidth. When the specified total bandwidth for the interface is reached, no further packets are sent, and any further packets may be dropped if the interface queues are full.

Packets are assigned to classes on the outbound interface based on either a mark assigned to the packet, or the ToS (type of service) field in the IP header. If the ToS field matches a defined class, then the packet is allocated to that class. Otherwise, it is allocated to any class that matches the mark

assigned to the packet, and if no class matches the mark, then the packet is assigned to the default class.

Marks are assigned to packets either by the TC Rules based on any of a number of parameters, such as IP address, port number, protocol, packet length, and so on.

39.1.2.1. Traffic Control Example

The goal of this example is to operate T1 port at 1.5Mbit/s and ensure that UDP source port 20000 traffic gets at least half the bandwidth, while ICMP and TCP ACK packets should have high priority, HTTP traffic gets at least 20% and at most 50%, and all other traffic should get what is left over but only up to 50% of the bandwidth.

The three TC menus would be configured as follows:

39.1.2.1.1. TC Interfaces

Interface	Inbound bandwidth	Outbound bandwidth
te1-3-1c24ppp	1500kbit	1500kbit

Table 39.1. TC Interfaces

39.1.2.1.2. TC Classes

Interface	Mark	Minimum	Maximum	Priority	Options
te1-3-1c24ppp	1	full/2	full	0	
te1-3-1c24ppp	2	28bps	full	1	tcp-ack
te1-3-1c24ppp	3	full/5	full*5/10	2	
te1-3-1c24ppp	4	28bps	full*5/10	3	default

Table 39.2. TC Classes

39.1.2.1.3. TC Rules

Mark	Source	Destination	Protocol	Source Port	Dest Port	Test	Length	TOS
2	Any	Any	ICMP	Any	Any	Any	Any	Any
RESTORE	Any	Any	Any	Any	Any	0	Any	Any
CONTINUE	Any	Any	Any	Any	Any	!0	Any	Any
1	Any	Any	UDP	20000	Any	Any	Any	Any
3	Any	Any	TCP	Any	80	Any	Any	Any
4	Any	Any	Any	Any	Any	0	Any	Any
SAVE	Any	Any	Any	Any	Any	!0	Any	Any

Table 39.3. TC Rules

The rules first check non connection-based protocol rules (ICMP in this case) in order to assign a mark. For any packet that is still not marked, we attempt to restore a saved mark for the connection. If at this point the packet has a mark set, we stop checking rules (CONTINUE) since it is either ICMP or a packet from an existing connection which we have already assigned a mark. If still no mark is assigned, it must be a new connection so we process the packet through all the remaining rules to determine the mark it should receive. At the end, we save the new mark to the connection so that any further packets for the connection do not have to go through all the rules again, in order to save processing resources. We mark all packets with no other matching rule to 4 since that represents the default class (as defined in TC Classes). This allows explicit traffic control of even unspecified network connections.

39.2. Traffic Control Configuration

Figure 39.1. Traffic-Control menu

To display the Traffic Control menu, navigate to *qos/traffic-control*.

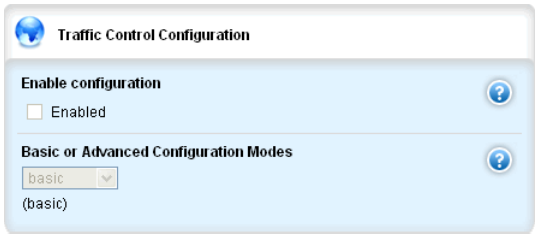

Figure 39.2. Traffic Control Configuration form

The Traffic Control Configuration form appears on the same screen as the Traffic Control menu.

Enable configuration

Enables/disables traffic control (TC) for the current firewall configuration. The current firewall configuration is the one that is committed. When an active configuration is committed to the system, then an **enabled** TC configuration will be included. When a work configuration is comitted then the **enabled** TC configuration will be included in the work config. **A TC configuration needs a firewall configuration to operate**.

Basic or Advanced Configuration Modes

Synopsis: *string - one of the following keywords { advanced, basic }*

Default: *basic*

Specifies to use either 'simple' or 'advanced' configuration modes. Click again on traffic-control after making a choice.

39.2.1. Traffic Control Modes

Traffic Control functions are divided into two modes: basic-configuration mode and advanced-configuration mode. Basic-configuration mode contains functions with basic traffic control configuration parameters. Advanced-configuration mode contains functions with advanced traffic control configuration parameters. The two modes cannot be accessed simultaneously. Only the mode that is currently configured can be accessed.

It is mandatory to configure the firewall before enabling traffic control. For information on configuring the firewall, refer to Chapter 38, Firewall.

39.2.1.1. Basic-configuration Mode

To configure basic-configuration mode, follow the procedure below.

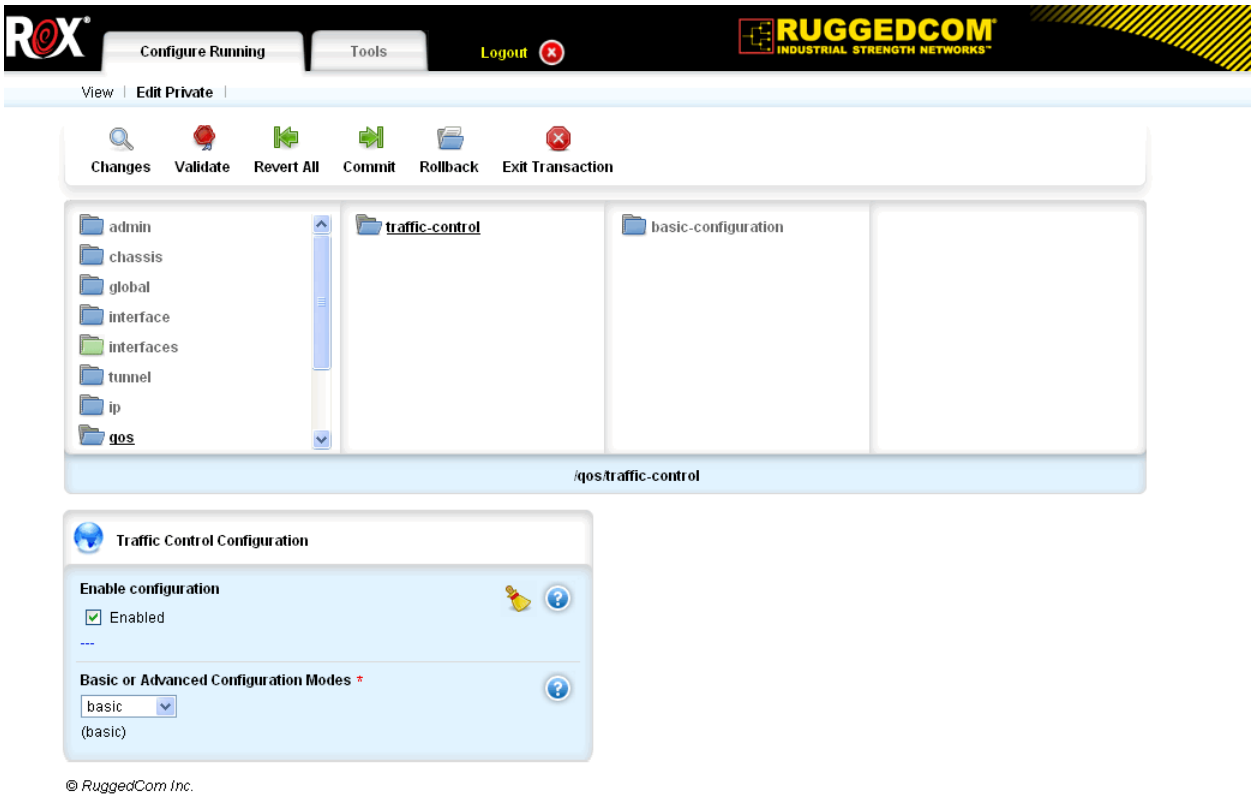

Figure 39.3. Enabling Basic-configuration Mode

Procedure 39.1. Configuring Basic-configuration Mode

1. Enter Edit Private mode.
2. Click on *qos/traffic-control*.
3. On the Traffic Control Configuration form, click **Enabled** in the *Enable configuration* field.
4. Select **basic** in the *Basic or Advanced Configuration Modes* field.
5. Click **Commit**.
6. Click **Exit Transaction**.

39.2.1.1.1. Interfaces

The interface is the network interface to which traffic shaping will apply.

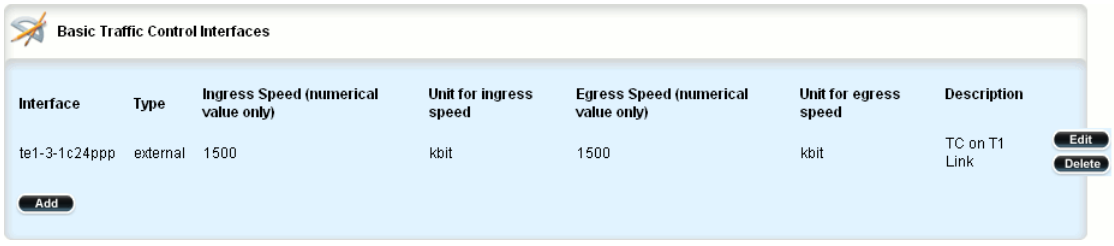

Figure 39.4. Basic Traffic Control Interfaces table

To display this table, navigate to *qos/traffic-control/basic-configuration/tcinterfaces*.

Figure 39.5. Interface to Apply Traffic Control form

To display this form, navigate to *qos/traffic-control/basic-configuration/tcinterfaces/{interface}*.

interface

Synopsis: *string*

An interface to which traffic shaping will apply

Type

Synopsis: *string - one of the following keywords { none, external, internal }*

Default: *none*

(optional) 'external' (facing toward the Internet) or 'internal'
(facing toward a local network). 'external'
causes the traffic generated by each unique
source IP address to be treated as a single
flow. 'internal' causes the traffic generated by
each unique destination IP address to be treated
as a single flow. , internal interfaces seldom
benefit from simple traffic shaping. Default: none.

Ingress Speed (numerical value only)

Synopsis: *string*

(optional) The incoming bandwidth of this interface. If incoming
traffic exceeds the given rate, received packets
are dropped randomly. When unspecified, maximum
speed is assumed. Specify only the number here.
The unit (kbps, mbps) is specified in In-unit.

Unit for ingress speed

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Specifies the unit when an incoming bandwidth is specified

Egress Speed (numerical value only)

Synopsis: *unsigned short integer*

The outgoing bandwidth for this interface. Specify only the number here. The unit (kbps, mbps) is specified in Out-unit.

Unit for egress speed

Synopsis: *string - one of the following keywords { bps, mbps, mbit, kbps, kbit }*

Specifies the unit for the outgoing bandwidth

Description

Synopsis: *string*

A description for this configuration item

> *Incoming traffic is prioritized based on the matching ToS value in the IP header if nothing else is configured under a band or when IP traffic does not match with the rules specified in a band.*

> *Speed units:*
> - *bps = bytes per second*
> - *mbps/kbps = megabyte/kilobyte per second*
> - *mbits/kbits = megabits/kilobits per second*

> *The Egress Speed and Unit for egress speed must be configured before committing a configuration.*

39.2.1.1.2. Priorities

Priorities configure traffic priorities for basic traffic shaping.

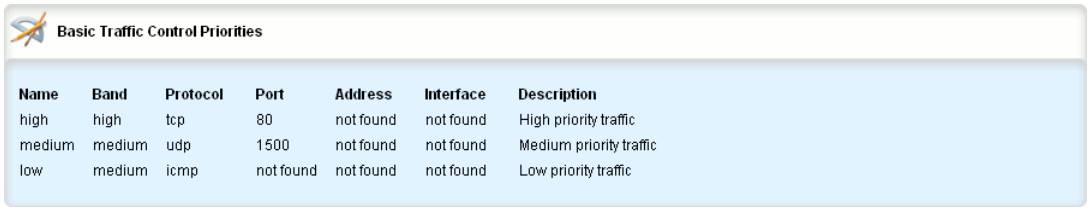
Basic Traffic Control Priorities
Name Band Protocol Port Address Interface Description
high high tcp 80 not found not found High priority traffic
medium medium udp 1500 not found not found Medium priority traffic
low medium icmp not found not found not found Low priority traffic

Figure 39.6. Basic Traffic Control Priorities table

To display this table, navigate to *qos/traffic-control/basic-configuration/tcpriorities*.

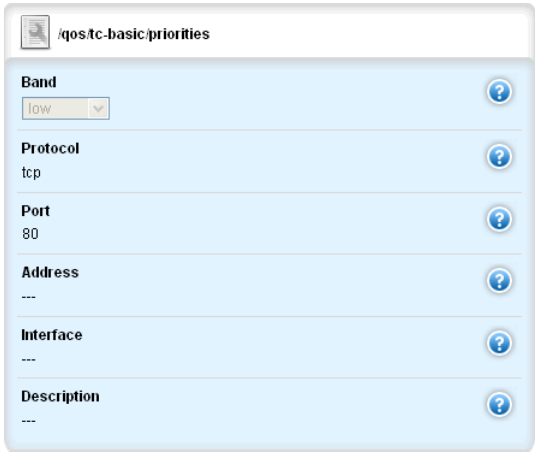

Figure 39.7. Priorities form

To display this form, navigate to *qos/traffic-control/basic-configuration/tcpriorities/{priority}*.

name

Synopsis: *string*

A distinct name for this configuration entry

band

Synopsis: *string - one of the following keywords { low, medium, high }*

Default: *medium*

Priority (band) : high, medium, low...

High band includes:

Minimize Delay (md) (0x10),

md + Minimize Monetary Cost (mmc) (0x12),

md + Maximize Reliability (mr) (0x14),

mmc+md+mr (0x16).

Medium band includes:

Normal Service (0x0),

mr (0x04),

mmc+mr (0x06),

md + Maximize Throughput (mt) (0x18),

mmc+mt+md (0x1a),

mr+mt+md (0x1c),

mmc+mr+mt+md (0x1e).

Low band includes:

mmc (0x02),

mt (0x08),

mmc+mt (0x0a),

mr+mt (0x0c),

mmc+mr+mt (0x0e).

protocol

Synopsis: *string*

Synopsis: *string - one of the following keywords { all, icmp, udp, tcp }*

(choice) Targeted protocol

port

Synopsis: *string*

(choice) Source port - can be specified **only if** protocol is tcp, udp, dccp, sctp or udplite

address

Synopsis: *string*

(choice) Source address - can be specified **only if** protocol, port and interface are not defined

interface

Synopsis: *string*

(choice) Source interface - can be specified **only if** protocol, port and address are not defined

description

Synopsis: *string*

(optional) A description for this configuration

> *For basic traffic control configurations, **Port**, **Address** and **Interface** refer to the source of the traffic.*

39.2.1.2. Advanced-configuration Mode

To configure advanced-configuration mode, follow the procedure below.

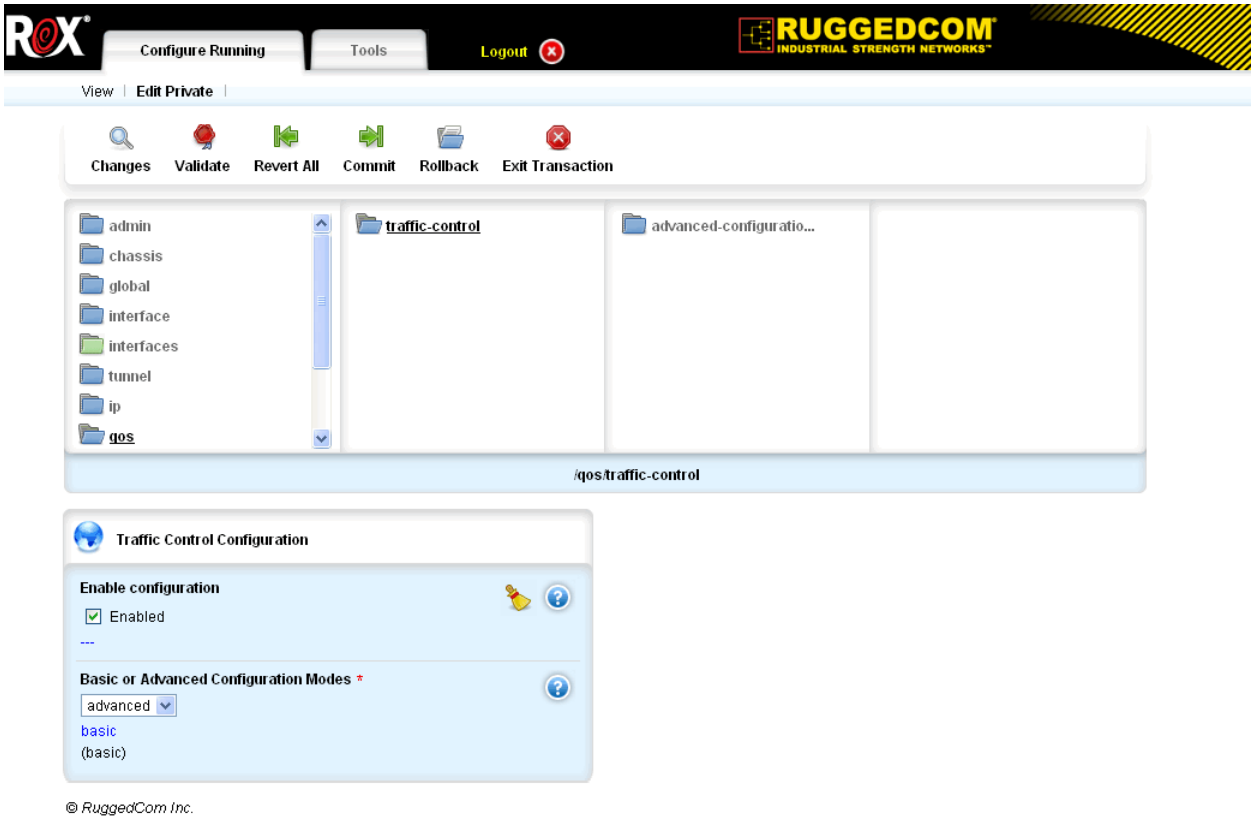

Figure 39.8. Enabling Advanced-configuration Mode

Procedure 39.2. Configuring Advanced-configuration Mode

1. Enter Edit Private mode.
2. Click on *qos/traffic-control*.
3. On the Traffic Control Configuration form, click **Enabled** in the *Enable configuration* field.
4. Select **advanced** in the *Basic or Advanced Configuration Modes* field.
5. Click **Commit**.
6. Click **Exit Transaction**.

39.2.1.2.1. Classes

Classes define classes for traffic shaping.

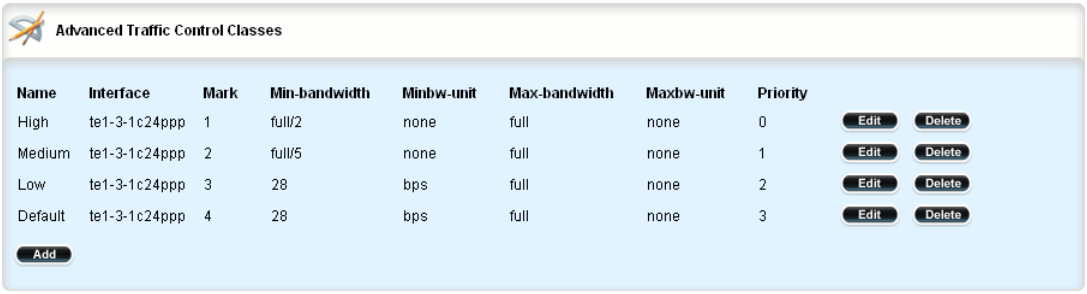
Advanced Traffic Control Classes
Name Interface Mark Min-bandwidth Minbw-unit Max-bandwidth Maxbw-unit Priority
High te1-3-1c24ppp 1 full/2 none full none 0 Edit Delete
Medium te1-3-1c24ppp 2 full/5 none full none 1 Edit Delete
Low te1-3-1c24ppp 3 28 bps full none 2 Edit Delete
Default te1-3-1c24ppp 4 28 bps full none 3 Edit Delete
Add

Figure 39.9. Advanced Traffic Control Classes table

To display this table, navigate to *qos/traffic-control/advanced-configuration/tcclasses*.

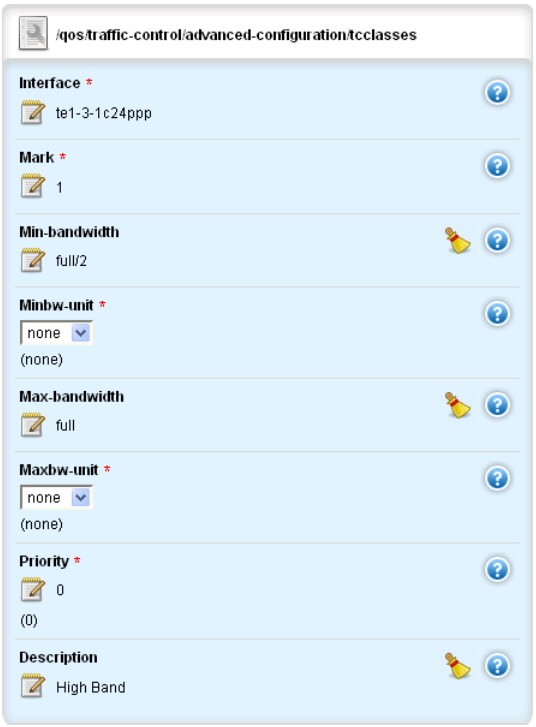

Figure 39.10. TC Classes form

To display this form, navigate to *qos/traffic-control/advanced-configuration/tcclasses/{class}*.

Note that each class is associated with exactly one network interface. Exactly one class for each interface must be designated as the default.

name

Synopsis: *string*

A name for this TC class entry

interface

Synopsis: *string*

The interface to which this class applies... Each interface must be listed only once.

mark

Synopsis: *unsigned short integer*

Mark that identifies traffic belonging to this class... This is an

unique integer between 1-255. Each class must have its own unique mark.

min-bandwidth

Synopsis: *string*

The minimum bandwidth this class should get,
when the traffic load rise... This can be
either a numeric value or a calculated expression
based on the bandwidth of the interface. A
fixed numerical value must only be a number -
its unit is specified in Minbw-unit.
A calculated expression is based on a fraction of the
'full' bandwidth, such as:
'full/3' for a third of the bandwidth and
'full*9/10' for nine tenths of the bandwidth.
In such a case, do not specify any Minbw-unit

minbw-unit

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Default: *none*

(kbps, mbps...) Only if min-bandwidth is a **single numerical value**

max-bandwidth

Synopsis: *string*

The maximum bandwidth this class is allowed to use
when the link is idle... This can be either a numeric
value or a calculated expression based on the bandwidth
of the interface. A fixed numerical value must only be
a number - its unit is specified in Maxbw-unit.
A calculated expression is based on a fraction of the
'full' bandwidth, such as:
'full/3' for a third of the bandwidth and
'full*9/10' for nine tenths of the bandwidth.
In such a case, do not specify any Maxbw-unit

maxbw-unit

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Default: *none*

(kbps, mbps...) Only if max-bandwidth is a **single numerical value**

priority

Synopsis: *unsigned short integer*

Default:

The priority in which classes will be serviced... Higher priority
classes will experience less delay since they are serviced
first. Priority values are serviced in ascending order
(e.g. 0 is higher priority than 1).

description

Synopsis: *string*

A description for this configuration item

Options

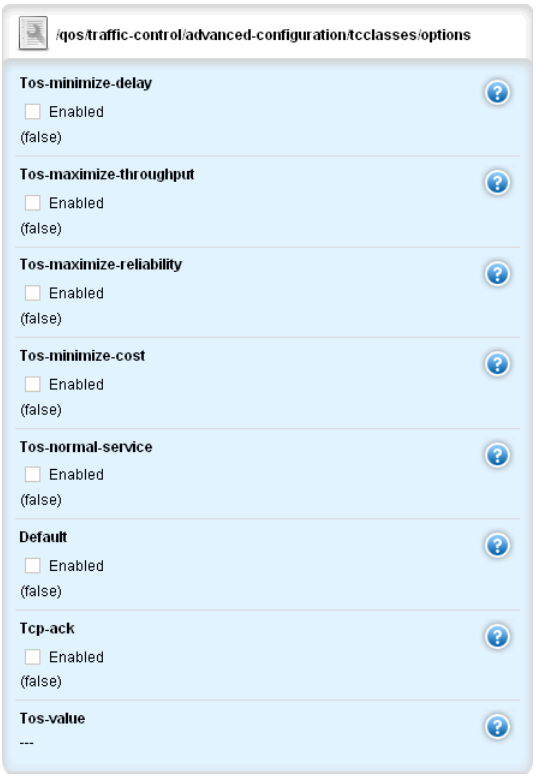

Figure 39.11. Options form

To display this form, navigate to *qos/traffic-control/advanced-configuration/tcclasses/{class}.*

IP Traffic matching with the ToS options take precedence over the mark rules.

tos-minimize-delay

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x10/0x10

tos-maximize-throughput

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x08/0x08

tos-maximize-reliability

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x04/0x04

tos-minimize-cost

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x02/0x02

tos-normal-service

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x00/0x1e

default

Synopsis: *boolean*

Default: *false*

One default class per interface **must** be defined

tcp-ack

Synopsis: *boolean*

Default: *false*

All tcp ack packets into this class... This option should be

specified only once per interface.

tos-value

Synopsis: *string*

Custom classifier for the given value/mask...

The values are hexadecimal, prefixed by '0x'.

Ex.:

0x56[/0x0F]

> *The ToS-value field allows you to define a classifier for the given ToS value or value/mask combination of an IP packet's TOS byte. Value and Value/Mask are both specified in hexadecimal notation using the "0x" prefix. It is also possible to specify a diffserv marking, or DSCP (Diffserv Code Point). These are typically quoted as 6-bit values, and must be left-shifted (multiplied by 4) for use in the "tos=" field. For example, a DSCP of 0x2E (EF, or Expedited Forwarding) would be entered as 0xB8/0xFC (4 X 0x2E = 0xB8, and the two lowest order bits are masked by masking with 0xFC).*

39.2.1.2.2. Devices

Devices characterize interfaces for traffic shaping, mostly for outbound bandwidth and the outgoing interface.

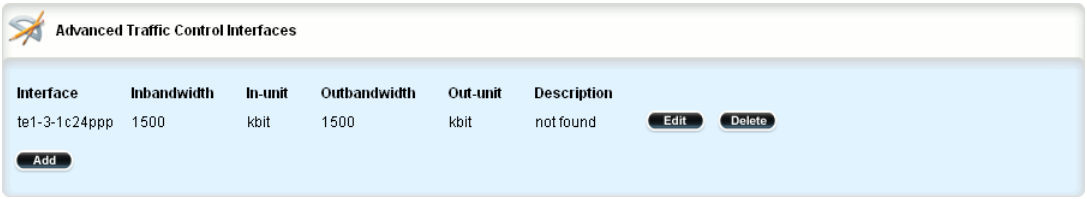

Figure 39.12. Advanced Traffic Control Interfaces table

To display this table, navigate to *qos/traffic-control/advanced-configuration/tcdevices*.

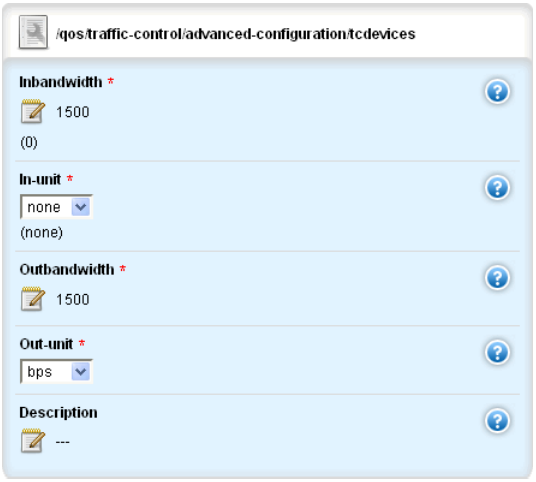

Figure 39.13. TC Devices form

To display this form, navigate to *qos/traffic-control/advanced-configuration/tcdevices/{interface}*.

interface

Synopsis: *string*

An interface to which traffic shaping will apply

inbandwidth

Synopsis: *unsigned short integer*

Default:

Incoming bandwidth - default: 0 = ignore ingress...

Defines the maximum traffic allowed for this interface in total,

if the rate is exceeded, the packets are dropped

in-unit

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Default: *none*

Unit when incoming bandwidth is specified

outbandwidth

Synopsis: *unsigned short integer*

Maximum outgoing bandwidth... This is the maximum speed that can be

handled. Additional packets will be dropped. This is the

bandwidth that can be refrred-to as 'full' when defining classes.

out-unit

Synopsis: *string - one of the following keywords { bps, mbps, mbit, kbps, kbit }*

Outgoing bandwidth unit

description

Synopsis: *string*

A description for this configuration item

39.2.1.2.3. Rules

Rules define packet marking rules, usually for traffic shaping.

Figure 39.14. TCrules menu

The *tcrules* menu allows you to add, edit or remove a traffic classification rule. Add a new rule by selecting <Add tcrules>. Remove a tcrule by selecting next to a tcrule and click on an existing tcrule to modify it. Reorder rules by clicking next to the rule submenus.

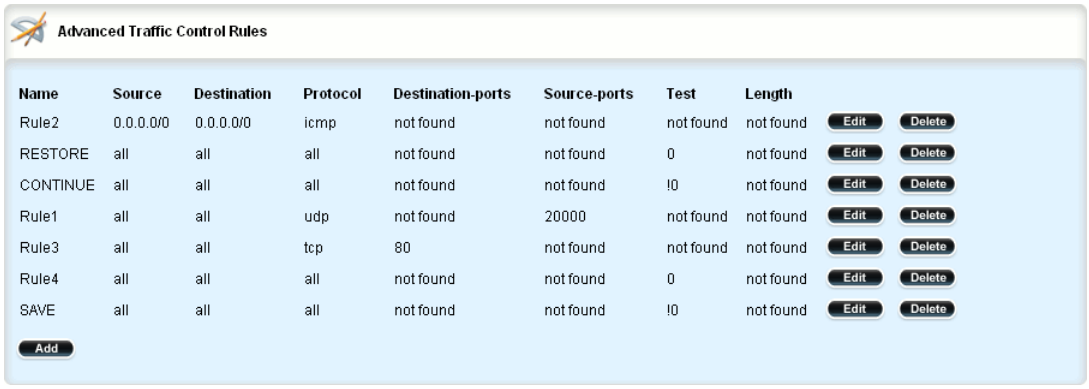
Advanced Traffic Control Rules
Name Source Destination Protocol Destination-ports Source-ports Test Length
Rule2 0.0.0.0/0 0.0.0.0/0 icmp not found not found not found not found Edit Delete
RESTORE all all all not found not found 0 not found Edit Delete
CONTINUE all all all not found not found !0 not found Edit Delete
Rule1 all all udp not found 20000 not found not found Edit Delete
Rule3 all all tcp 80 not found not found not found Edit Delete
Rule4 all all all not found not found 0 not found Edit Delete
SAVE all all all not found not found !0 not found Edit Delete
Add

Figure 39.15. Advanced Traffic Control Rules table

To display this table, navigate to *qos/traffic-control/advanced-configuration/tcrules*.

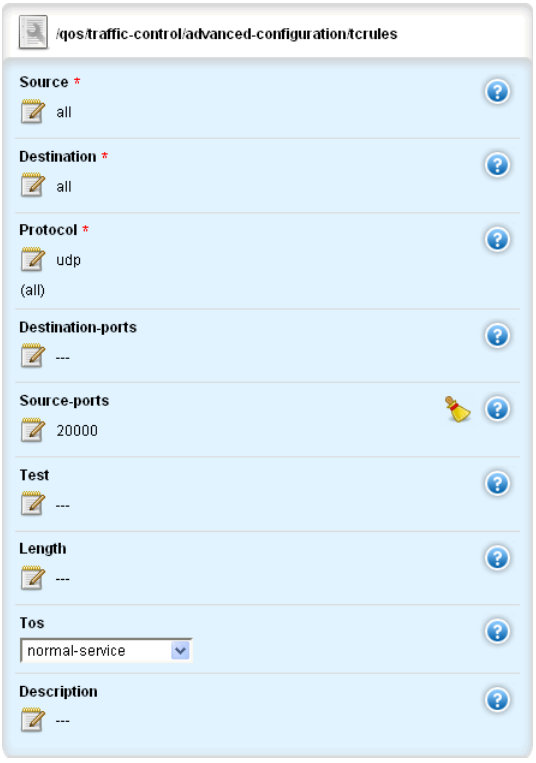

Figure 39.16. TCrules form

To display this form, navigate to *qos/traffic-control/advanced-configuration/tcrules/{rule}*.

name

Synopsis: *string*

A distinct name for this rule

source

Synopsis: *string*

IF name, comma-separated list of hosts or IPs, MAC addr, or 'all'...

When using MACs, use '~' as prefix and '-' as separator. Ex.:

~00-1a-6b-4a-72-34,~00-1a-6b-4a-71-42

destination

Synopsis: *string*

IF name, comma-separated list of hosts or IPs, or 'all'

protocol

Synopsis: *string*

Synopsis: *string - one of the following keywords { all, icmp, udp, tcp }*

Default: *all*

The protocol to match - default: all

destination-ports

Synopsis: *string*

(Optional) Comma-separated list of port names, port numbers or port ranges

source-ports

Synopsis: *string*

(Optional) Comma- separated list of port names, port numbers or port ranges

test

Synopsis: *string*

(Optional) Defines a test on the existing packet or connection mark...

Default is packet mark. For testing a connection mark, add

':C' at the end of the test value. Ex.:

Test if the packet mark is not zero:

!0

Test if the connection mark is not zero:

!0:C

length

Synopsis: *string*

(Optional) Match the length of a packet against a specific value or range of values...

Greater than and lesser than, as well as ranges are supported in the form of min:max. Ex.:

Equal to 64

64

Greater or equal to 65

65:

Lesser or equal to 65

:65

In-between 64 and 768

64:768

tos

Synopsis: *string*

Synopsis: *string - one of the following keywords { normal-service, minimize-cost, maximize-reliability, maximize-throughput, minimize-delay }*

(Optional) Type Of Service ...

A pre-defined TOS value or a numerical value. The

numerical value is hexadecimal. Ex.: 0x38

description

Synopsis: *string*

A description for this configuration item

Mark-choice

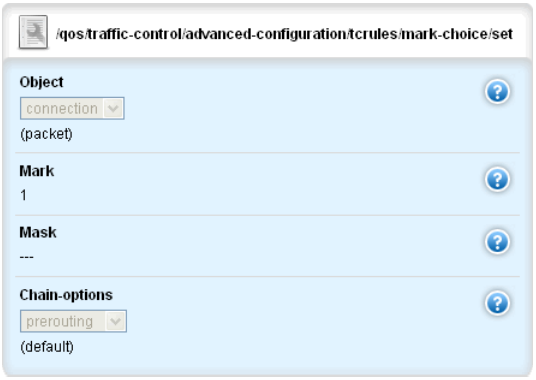

Figure 39.17. Set form

object

Synopsis: *string - one of the following keywords { connection, packet }*

Default: *packet*

Set the mark on either a packets or a connection

mark

Synopsis: *string*

Mark that corresponds to a class mark (decimal value)

mask

Synopsis: *string*

(optional) Mask to determine which mark bits will be set

chain-options

Synopsis: *string - one of the following keywords { prerouting, postrouting, forward }*

Default: *forward*

Chain where the set operation will take place

The chain-options field specifies the chain in which the rule will be processed.

- *Prerouting - Mark the connection in the PREROUTING chain.*

  This can be used with DNAT, SNAT and Masquerading rule in firewall. An example of such a rule is "Source.IP:192.168.2.101, Chain-option: preroute or default" but the actual Source.NAT address is 2.2.2.2.

- *Postrouting - Mark the connection in the POSTROUTING chain.*

  This can be used with DNAT, SNAT and Masquerading rules in the firewall. An example of such rule is "Destination.IP:192.168.3.101, Chain-option:preroute or default". In this case, the actual destination address is 192.168.3.101 but it will be translated to 192.168.3.33 by DNAT. Another example of a traffic control rule is ""Destination.IP:192.168.3.33, Chain-option:postrouting".

- *Forward - Mark the connection in the FORWARD chain.*

  This is the default chain option and it can be used for normal IP traffic without any address or port translation.

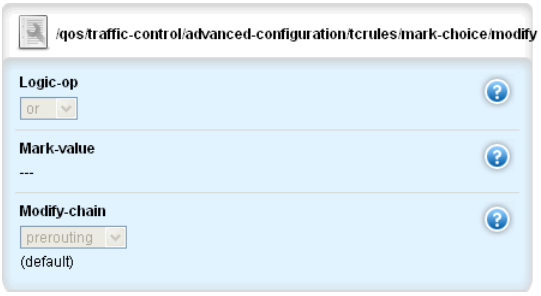

Figure 39.18. Modify form

logic-op

Synopsis: *string - one of the following keywords { or, and }*

Logical operation to perform on the current mark: AND/OR

mark-value

Synopsis: *string*

Mark to perform the operation with (decimal value)

modify-chain

Synopsis: *string - one of the following keywords { prerouting, postrouting, forward }*

Default: *forward*

Chain in which the operation will take place

Figure 39.19. Save form

value-mask

Synopsis: *string*

Mask to process the mark with

op-chain

Synopsis: *string - one of the following keywords { prerouting, forward }*

Default: *forward*

Chain in which the operation will take place

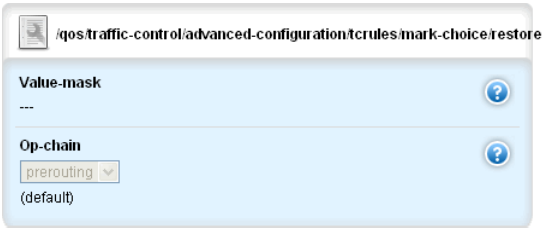

Figure 39.20. Restore form

value-mask

Synopsis: *string*

Mask to process the mark with

op-chain

Synopsis: *string - one of the following keywords { prerouting, forward }*

Default: *forward*

Chain in which the operation will take place

Figure 39.21. Continue form

continue-chain

Synopsis: *string - one of the following keywords { prerouting, forward }*

Default: *forward*

Chain in which the operation will take place

Hints on optimizing the TC Rule table

Every rule is processed in table order for every packet, unless a CONTINUE rule is matched, in which case processing stops. This can be used to improve efficiency in combination with the SAVE and RESTORE rules. For example, consider a TC Rules table organized roughly as follows (and in the same order):

- A RESTORE rule is used to restore the connection's mark to a matching, unmarked packet
- A CONTINUE if the mark is non zero
- Specific rules to check criteria to assign a mark, and finally
- A SAVE mark to connection if the mark is non zero (that is, a match was found above)

Using the above structure for the TC Rules table, only the first packet of any tcp or udp connection will have to go through all the rules, while every following packet will have its mark restored by the first rule, and then CONTINUE, skipping potentially many matching rules in the remainder of the table.

40. VRRP

40.1. VRRP Fundamentals

The Virtual Router Redundancy Protocol (VRRP) eliminates a single point of failure associated with statically routed networks by providing automatic failover using alternate routers. The RuggedBackbone™ VRRP daemon (keepalived) is an RFC 2338 version 2 compliant implementation of VRRP.

40.1.1. The Problem With Static Routing

Many network designs employ a statically configured default route in the network hosts. A static default route is simple to configure, requires little if any overhead to run and is supported by virtually every IP implementation. When dynamic host configuration protocol (DHCP) is employed, hosts may accept configuration for only a single default gateway.

Unfortunately, this approach creates a single point of failure. Loss of the router supplying the default route or the router's WAN connection results in isolating the hosts relying upon the default route.

There are a number of ways that may be used to provide redundant connections to the host. Some hosts can configure alternate gateways while others are intelligent enough to participate in dynamic routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First routing protocol (OSPF). Even when available, these approaches are not always practical due to administrative and operation overhead.

40.1.2. The VRRP Solution

VRRP solves the problem by allowing the establishment of a "virtual router group", composed of a number of routers that provide a specific default route. VRRP uses an election protocol to dynamically assign responsibility for the "virtual" router to one of the routers in the group. This router is called the VRRP Master. If the Master (or optionally its WAN connection) fails, the alternate (i.e. backup) routers in the group elect a new Master. The new master provides the virtual IP address and issues a gratuitous ARP to inform the network of where the gateway can be reached.

Because the host's default route does not change and MAC address is updated, packet loss at the hosts is limited to the amount of time required to elect a new router.

40.1.3. VRRP Terminology

Each physical router running VRRP is known as a VRRP Router. Two or more VRRP Routers can be configured to form a "Virtual Router". Each VRRP Router may participate in one or more Virtual Routers.

Each Virtual Router has a user-configured *Virtual Router Identifier* (VRID) and an *Virtual IP address* or set of IP addresses on the shared LAN. Hosts on the shared LAN are configured to use these addresses as the default gateway.

One router in the Virtual Router Group will be elected as the *Master*, all other routers in the group will be *Backups*.

Each router in the group will run at a specific *Priority*. The router with the highest priority is elected Master. The value of Priority varies from 1 to 255.

VRRP can also monitor a specified interface and give up control of a VRIP if that interface goes down.

In the following network, host 1 uses a gateway of 1.1.1.253 and host 2 uses a gateway of 1.1.1.252. The 1.1.1.253 gateway is provided by VRID 10. In normal practice router 1 will provide this virtual IP as its priority for VRID 10 is higher than that of router 2. If router 1 becomes inoperative or if its w1ppp link fails, it will relinquish control of VRIP 1.1.1.253 to router 2.

In a similar fashion host 2 can use the VRID 11 gateway address of 1.1.1.252 which will normally be supplied by router 2.

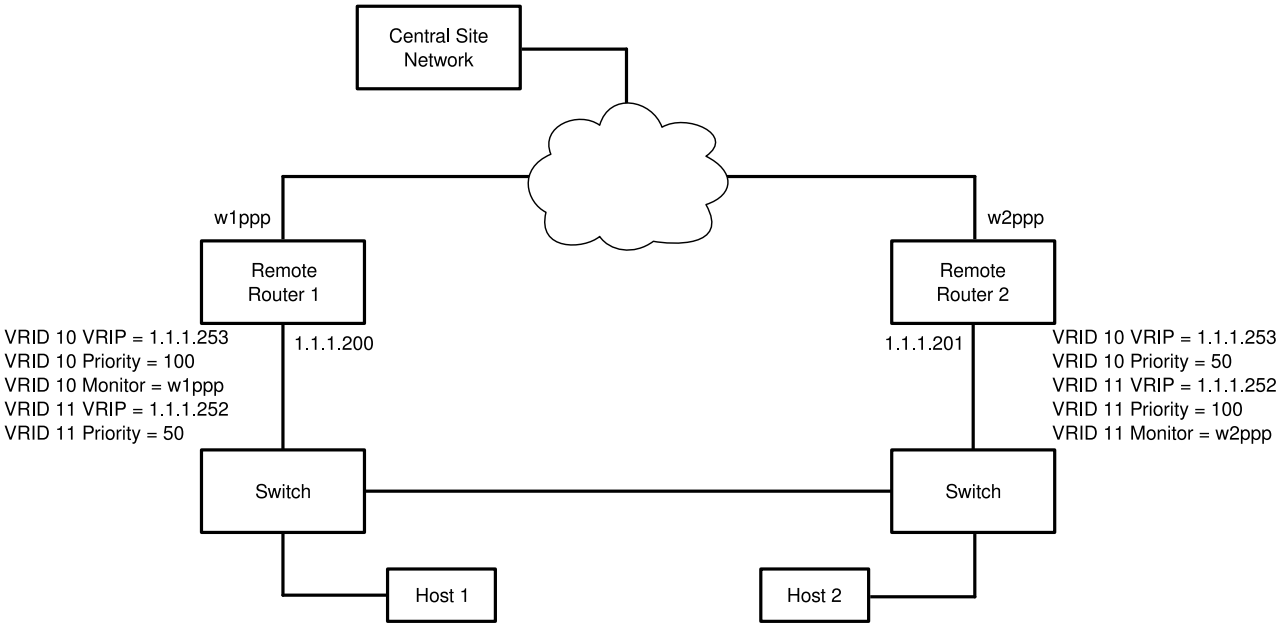

Figure 40.1. VRRP Example

In this example traffic from host1 will be sent through router 1 and traffic from host2 through router 2. A failure of either router (or its wan link) will be recovered by the other router.

Note that both routers can always be reached by the hosts at their “real” IP addresses.

Two or more VRRP instances can be assigned to be in the same *VRRP Group*, in which case, they can fail over together.

In the following network, both host 1 and host 2 use a gateway of 192.168.3.10. The external side can access the internal side by gateway 192.168.2.10. The VRID_20 and VRID_21 are grouped together. Normally the Router 1 will provide both internal and external access gateway as its priority is higher than those on Router 2. When either internal or external side of Router 1 becomes inoperative, it will remove the control of both VRIP 192.168.2.10 and 192.168.3.10 and gives the control to Router 2.

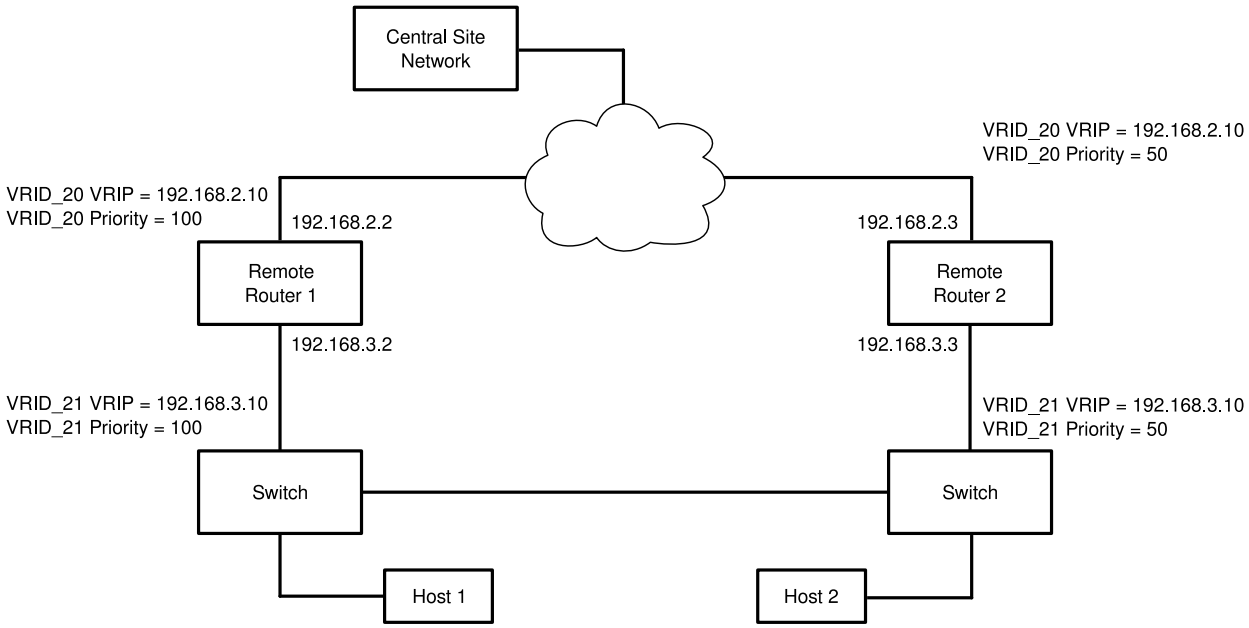

Figure 40.2. VRRP Group Example

Other VRRP parameters are the *Advertisement Interval* and *Gratuitous ARP* Delay.

The advertisement interval is the time between which advertisements are sent. A backup router will assume mastership *four advertisement intervals* after the master fails, so the minimum fail-over time is four seconds. If a monitored interface goes down, a master router will immediately signal an election and allow a backup router to assume mastership.

The router issues a set of gratuitous ARPs when moving between master and backup state. These unsolicited ARPs teach the hosts and switches in the network of the current MAC address and port associated with the VRIP. The router will issue a second set of ARPs after the time specified by the Gratuitous ARP delay.

40.2. VRRP Configuration

Figure 40.3. VRRP Menu

The VRRP menu is accessible from the main menu under *services* at *services/vrrp*.

Figure 40.4. Virtual Router Redundancy Protocol (VRRP) Form

The Virtual Router Redundancy Protocol (VRRP) form appears on the same screen as the *VRRP* menu. In the Virtual Router Redundancy Protocol (VRRP) form, enable or disable the VRRP service.

Enable VRRP Service

Enables VRRP Service.

Router ID

Synopsis: *string*

The router ID for VRRP logs.

Figure 40.5. VRRP Group Table

The VRRP Group table displays configured VRRP groups. To display this table, navigate to *services/vrrp/group*.

Group Name

Synopsis: *string*

The group name.

Figure 40.6. VRRP Instance Table

To display this table, navigate to *services/vrrp/instance*.

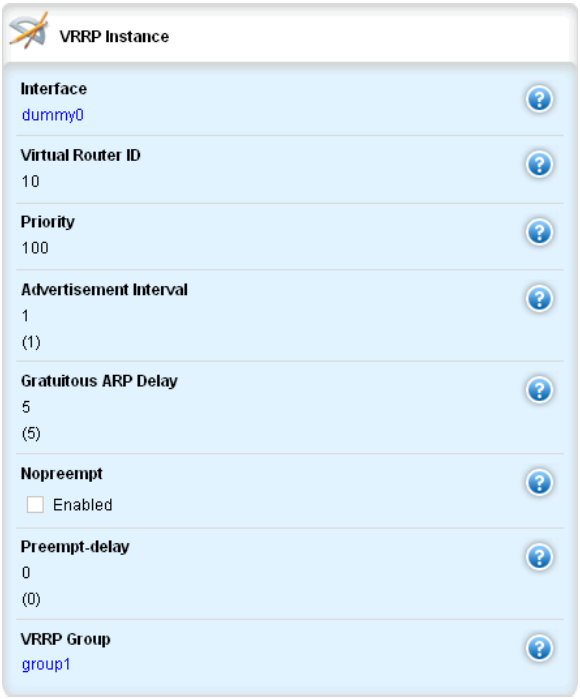

Figure 40.7. VRRP Instance Form

The VRRP Instance Form is used when configuring a VRRP instance. To display this form, navigate to *services/vrrp/instance/VRID20*.

Instance Name

Synopsis: *string*

The VRRP instance name.

Interface

Synopsis: *A string*

The interface that VRRP packets are sent on.

Virtual Router ID

Synopsis: *unsigned byte*

The Virtual Router ID. All routers supplying the same VRIP should have the same VRID.

Priority

Synopsis: *unsigned byte*

The priority of VRRP instance. For electing MASTER, highest priority wins.

Advertisement Interval

Synopsis: *unsigned byte*

Default: *1*

The advertisement interval (in seconds).

Gratuitous ARP Delay

Synopsis: *unsigned byte*

Default: *5*

Gratuitous ARP delay (in seconds). This controls the number of seconds after the router changes between the master and the backup state before a second set of gratuitous ARPs are sent.

nopreempt

Allows lower priority machine to maintain master role, even when a higher priority machine comes back online.

preempt-delay

Synopsis: *unsigned integer*

Default:

Seconds after startup until preemption.

use-virtual-mac

Use virtual MAC.

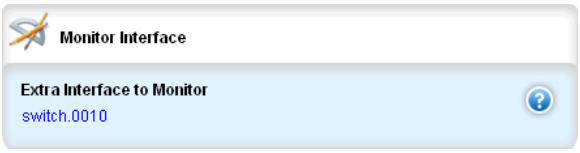

Figure 40.8. Monitor Interface Form

To display this form, navigate to *services/vrrp/instance/VRID20/monitor*.

An Extra Interface to Monitor causes VRRP to release control of the VRIP if the specified interface stops running.

Extra Interface to Monitor

Synopsis: *A string*

The interface name.

Figure 40.9. VRIP IP Address Table

The VRIP IP Address Table shows configured VRIP IP addresses associated with a VRID. To display this table, navigate to *services/vrrp/instance/VRID20/vrip*.

Virtual IP Address/Netmask

Synopsis: *IPv4 address and prefix in CIDR notation*

Virtual IP Address/Netmask.

40.2.1. VRRP Status

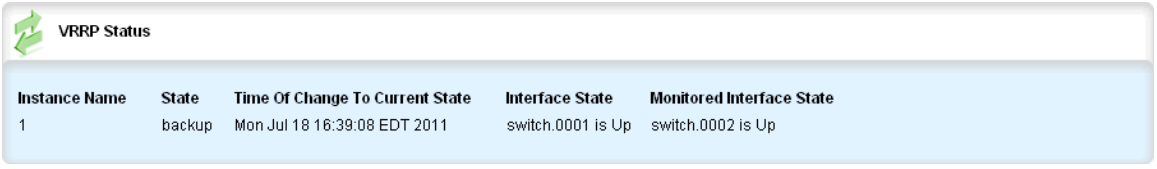

Figure 40.10. VRRP Status Table

To display this table, navigate to *services/vrrp/status*.

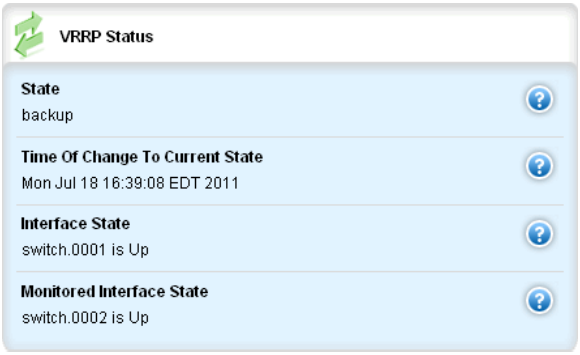

Figure 40.11. VRRP Status Form

To display this form, navigate to *services/vrrp/status/{number}*.

Instance Name

Synopsis: *string*

The VRRP instance name.

State

Synopsis: *string*

The VRRP instance state.

Time Of Change To Current State

Synopsis: *string*

The time of change to the current state.

Interface State

Synopsis: *string*

The VRRP interface state.

Monitored Interface State

Synopsis: *string*

Monitors the interface state.

41. Link Failover

Link failover provides an easily configured means of raising a backup link upon the failure of a designated main link. The main and backup links can be Ethernet, Cellular Modem, T1/E1, or DDS.

Link failover can back up to multiple remote locations, managing multiple main-to-backup link relationships. When the backup link is a modem, many “profiles” of dialed numbers can exist, each serving as a distinct backup link.

Link failover can back up a permanent, high-speed WAN link to a permanent, low-speed WAN link. Use this function when OSPF cannot be employed, such as on public links.

You can also use link failover to migrate the default route from the main link to the backup link.

The time after a main link failure to backup link startup, and the time after a main link recovery to backup link stoppage, are configurable. The link failover function also provides failover status information and a test of the failover settings.

41.1. Path Failure Discovery

To discover a primary path failure (in this example, through Network A), the link backup daemon inspects the link status of the main link and sends a regular ping to a designated host (or to a dummy address on the router). In this way, network link failures can be discovered. It is essential that the designated host or that the dummy address always respond to the ping.

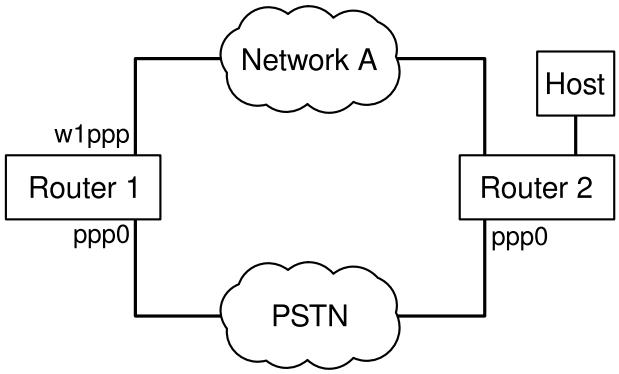

Figure 41.1. Link Backup Example

The daemon considers the main link to have failed, even if its link status is “up”, if the remote host fails to respond to a configurable number of pings after waiting a configurable timeout for each ping.

41.2. Using Routing Protocols and the Default Route

If the main trunk is on a private network, use a routing protocol to ensure that an alternate route to the end network is learned after the backup trunk is raised. Ensure that OSPF/RIP are configured to operate on the secondary trunk, assigning the secondary trunk a higher metric cost than that of the main trunk.

If the main trunk is on a public network, use the “transfer default route” feature.

The default route of the main trunk to be backed up using “transfer default route” must be assigned a metric greater than one.

41.3. Configuring Link Failover

To view the list of configured link failover settings, navigate to */services/link-failover*.

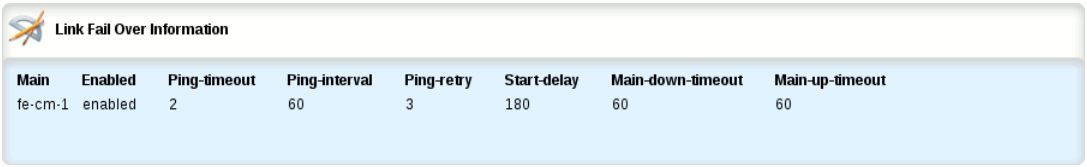

Figure 41.2. Link Fail Over Information Table

To configure link failover, do the following:

- set the link failover settings. See Section 41.3.1, “Configuring the Link Failover Settings”.
- add a link failover backup interface. See Section 41.3.2, “Setting a Link Failover Backup Interface”.
- add a link failover ping target. See Section 41.3.3, “Setting a Link Failover Ping Target”.
- enable link backup on demand. See Section 41.3.4, “Link Backup On Demand”.

After configuring link failover, you can do the following:

- view the link failover status. See Section 41.3.5, “Viewing Link Failover Status”.
- view the link failover log. See Section 41.3.6, “Viewing the Link Failover Log”.
- test the link failover settings. See Section 41.3.7, “Testing Link Failover”.

41.3.1. Configuring the Link Failover Settings

To configure the link failover settings:

- In edit mode, navigate to */services/link-failover* and click **<Add link-failover>**.
- On the *Key settings* form, select an interface from the list and click **Add**.
- On the *Link Fail Over Settings* form, set the link failover parameters.
- Commit the changes.

Figure 41.3. Link Fail Over Settings form

enabled

Enables this link backup.

ping-timeout

Synopsis: *integer*

Default: *2*

The time interval, in seconds, before immediately retrying a ping.

ping-interval

Synopsis: *integer*

Default: *60*

The time interval, in seconds, between ping tests.

ping-retry

Synopsis: *integer*

Default: *3*

The number of ping retries before construing a path failure.

start-delay

Synopsis: *integer*

Default: *180*

The delay time, in seconds, when first starting link failover.

main-down-timeout

Synopsis: *integer*

Default: *60*

The delay time, in seconds, that the main trunk is down before starting the backup trunk.

main-up-timeout

Synopsis: *integer*

Default: *60*

The delay time, in seconds, to confirm that the main trunk is up (returned to service) before stopping the backup trunk.

The link failover feature can only be configured on a routable interface. For the link failover feature to be used on a switched port, another VLAN must be configured (for example, switch.0002) to logically differentiate the switched port from the default PVID VLAN 1 (switch.0001).

41.3.2. Setting a Link Failover Backup Interface

A backup interface is the interface to which link failover switches when the main interface is determined to be down. You can add up to three backup interfaces to each link failover configuration.

To set a link failover backup interface:

- In edit mode, navigate to */services/link-failover{link failover ID}/backup* and click **<Add backup>**.
- On the *Key settings* form, select an interface from the list and click **Add**.
- On the *Backup Settings* form, set the backup parameters.
- Commit the changes.

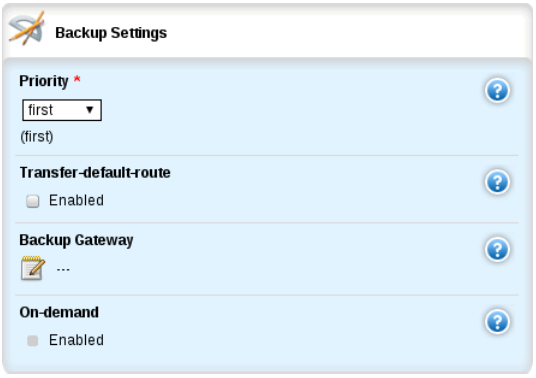

Figure 41.4. Backup Settings form

priority

Synopsis: *string - one of the following keywords { first, second, third }*

Default: *first*

The priority which is applied to the backup interface when switching

transfer-default-route

Transfer default gateway on switching main and backup interface.

The default route on the main interface must have a 'distance' greater than one.

Backup Gateway

Synopsis: *A string conforming to: "(([0-9]\|[1-9][0-9]\|1[0-9]{2}\|2[0-4][0-9]\|25[0-5])\.){3}([0-9]\|[1-9][0-9]\|1[0-9]{2}\|2[0-4][0-9]\|25[0-5])"*

The IP address of the backup gateway

on-demand

Synopsis: *boolean*

Displays the status of the interface's On-demand option. When enabled, link failover can set the interface up or down as needed; the interface is down until needed by link failover. When disabled, link failover cannot set the interface up or down; by default, the interface is always up.

41.3.3. Setting a Link Failover Ping Target

A link failover ping target is an IP address that link failover pings to determine if the main link is down. The address can be a dedicated host or a dummy address on a router. You can add up to three link failover ping targets to each link failover configuration.

To set a link failover ping target:

- In edit mode, navigate to */services/link-failover{link failover ID}/backup* and click **<Add target>**.
- On the *Key settings* form, enter the IP address of the host to ping and click **Add**.
- Commit the changes.

Link failover pings each target separately. If all targets are down, the main link is considered to be down and it fails over to the backup interface. Backup links are used in the order of their Priority setting (first, second, and then third), always starting with the first priority interface. When a higher-priority interface becomes available again, the system reverts to the higher priority interface. For example, if the second priority interface is active, the system switches back to the first priority interface when the first priority interface becomes available again.

41.3.4. Link Backup On Demand

Use the *On-demand* option to keep interfaces down until they are needed by link failover:

- When the *On-demand* option is enabled on an interface, the interface is down by default. The interface is brought up when needed by the link failover function, and is brought down again when no longer needed.
- When the *On-demand* option is not enabled on an interface, the interface is always up by default.

The *On-demand* option can be set for the following interfaces:

- eth: on the Routable Ethernet Ports form at */interface/eth{interface id}*
- wan hdlc-eth connections: on the HDLC-ETH form at */interface/wan{interface id}/t1/channel{channel id}/connection/hdlc-eth*
- wan mlppp connections: on the Multilink PPP form at */interface/wan{interface id}/t1/channel{channel id}/connection/mlppp*
- wan ppp connections: on the PPP form at */interface/wan{interface id}/t1/channel{channel id}/connection/ppp*
- wan framerelay connections: on the DLCI form at */interface/wan{interface id}/t1/channel{channel id}/connection/framerelay/dlci*

After enabling the on-demand option, the On-demand field indicates the option's status on the Backup Settings form.

> *The On-demand feature is not available on switched ports, even though the link failover function is available on VLAN switched ports. The On-demand feature is available only on WAN and ETH interfaces.*

41.3.5. Viewing Link Failover Status

The *Link Fail Over Status* form displays the current link failover status. To view the link failover status in normal or edit mode, navigate to */services/link-failover{interface id}/status*.

Figure 41.5. Link Fail Over Status form

main-link-status

Synopsis: *string*

The main link status.

backup-link-status

Synopsis: *string*

The backup link status.

main-ping-test

Synopsis: *string*

Results of the pinging target using the main interface.

time-of-last-state-change

Synopsis: *string*

The time of the last state change.

link-backup-state

Synopsis: *string*

The backup link state.

backup-interface-in-use

Synopsis: *string*

The name of the backup interface that is being used.

41.3.6. Viewing the Link Failover Log

The *Link Fail Over Logs* form displays a log of link failover events. To view the link failover log in normal or edit mode, navigate to */services/link-failover{interface id}/log* and click **Perform**.

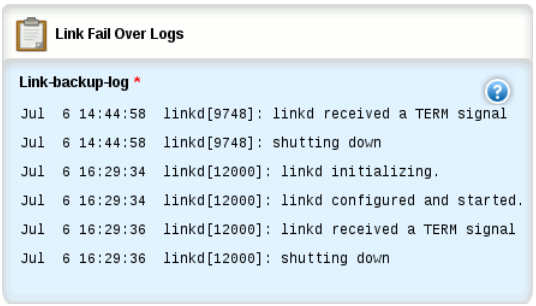
Link Fail Over Logs
Link-backup-log *
Jul 6 14:44:58 linkd[9748]: linkd received a TERM signal
Jul 6 14:44:58 linkd[9748]: shutting down
Jul 6 16:29:34 linkd[12000]: linkd initializing.
Jul 6 16:29:34 linkd[12000]: linkd configured and started.
Jul 6 16:29:36 linkd[12000]: linkd received a TERM signal
Jul 6 16:29:36 linkd[12000]: shutting down

Figure 41.6. Link Fail Over Logs form

41.3.7. Testing Link Failover

You can test your link failover settings to confirm that each link failover configuration works properly. To launch the test, you specify for how long the system should operate on the backup interface, and for how long the system should delay before starting the test. You can also cancel a test while it is in progress. Cancelling the test returns the interfaces to their pre-test condition.

While the test is running, monitor the *Link Fail Over Status* form to observe the main and backup link status, ping test results, state change, backup state, and backup interface information. As the test progresses, this information changes as link failover switches from the main interface to the backup interface. For more information on the *Link Fail Over Status* form, see Section 41.3.5, “Viewing Link Failover Status”.

To launch a link failover test:

- In normal mode or edit mode, navigate to */services/link-failover{interface id}/start-test*.
- On the *Link Fail Over Test Settings* form, set the test parameters.
- On the *Trigger Action* form, click **Perform**.

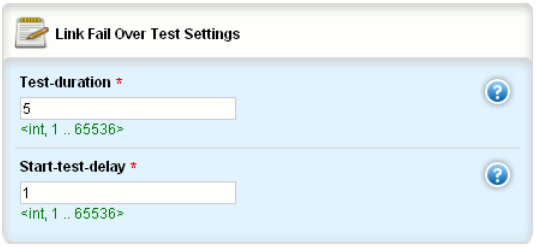

Figure 41.7. Link Fail Over Test Settings form

Test-duration

The amount of time (in minutes) to run the test before restoring service to the main trunk.

Start-test-delay

The amount of waiting time (in minutes) before running the test.

Part IV. Appendices

Upgrading Software	Appendix A, *Upgrading Software*
RADIUS Server Configuration	Appendix B, *RADIUS Server Configuration*
Setting Up An Upgrade Server	Appendix C, *Setting Up An Upgrade Server*
Adding and Replacing Modules	Appendix D, *Adding and Replacing Line Modules*
GNU General Public License	Appendix E, *GNU General Public License*

Appendix A. Upgrading Software

To launch a ROX™ operating system software upgrade, follow the procedure outlined below.

A.1. Preparing The Software Upgrade

The first step in a ROX™ software upgrade is to configure the location of the software upgrade repository and the version of software to which to upgrade.

At the top of the screen, click **Edit Private** to access the Edit Private view. The screen in Edit Private view is shown in the figure below.

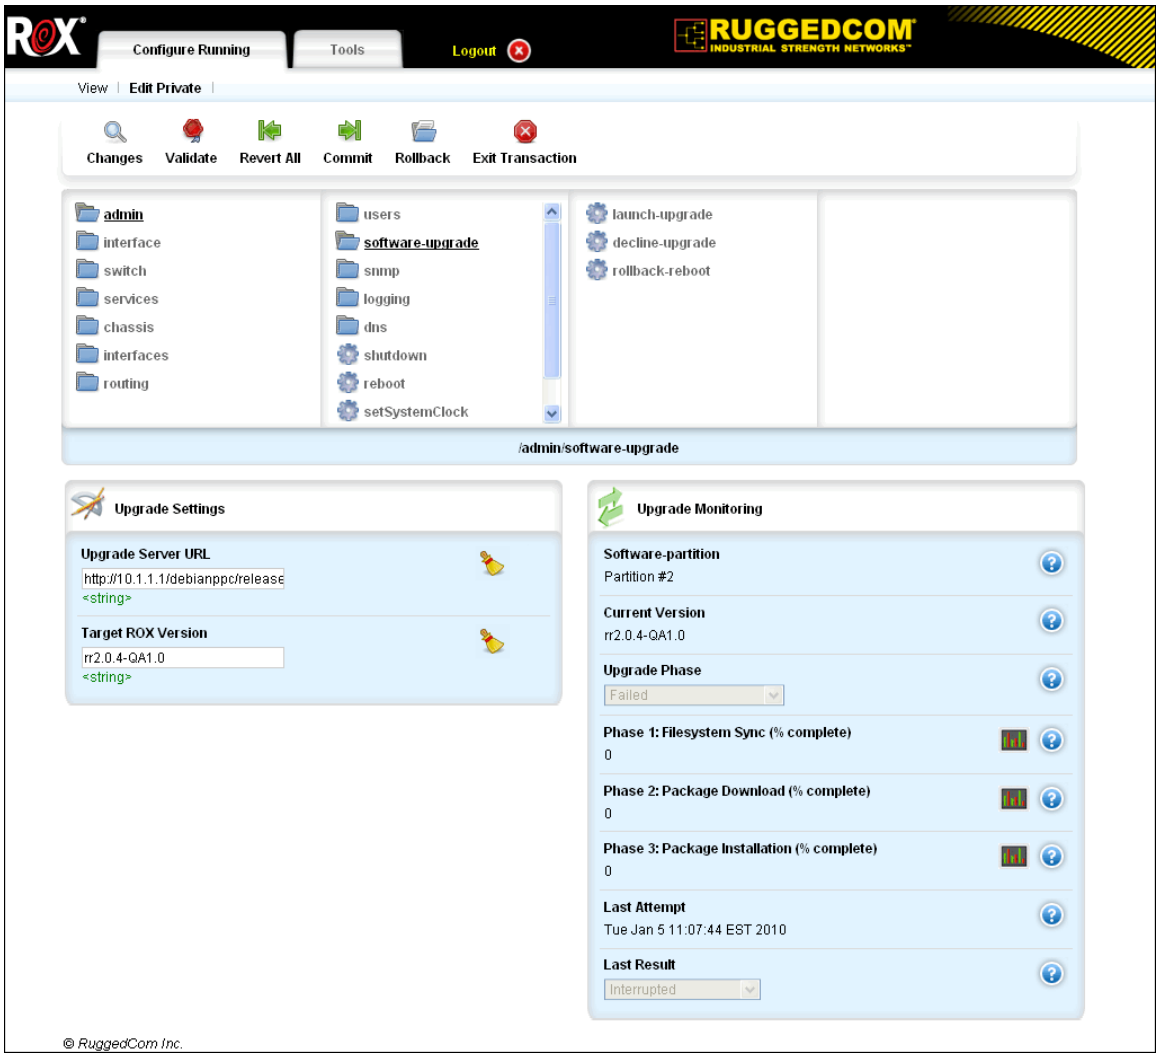

Figure A.1. The Software Upgrade Menu Interface

In the Upgrade Settings form, enter the location of the software upgrade server Upgrade Server URL field, and the version string of the desired version of ROX™ in the Target ROX Version field.

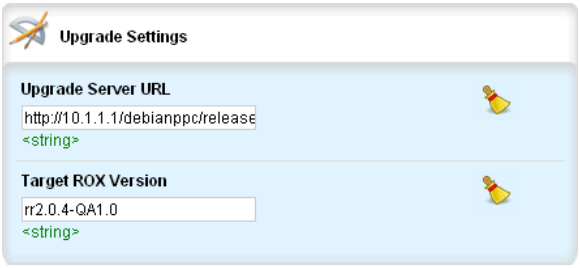

Figure A.2. Entry Fields in Upgrade Settings Form

After completing the information in the Upgrade Settings form, click the **Commit** button ( ) at the top of the screen. A dialog box will appear, prompting you to commit your changes. Click the **OK** button.

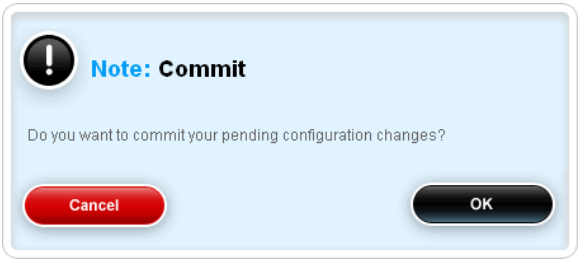

Figure A.3. Pending Commit

A dialog box will appear, informing you that the configuration has been committed. Click the **OK** button.

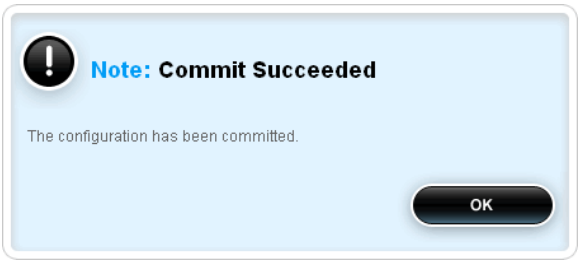

Figure A.4. Commit Succeeded

A.2. Launching The Upgrade

Click on the *launch-upgrade* menu action. Then, click the **Perform** button on the Launch Upgrade form. A dialog box will assert that: "Configurations will be locked until next boot." Click the **OK** button.

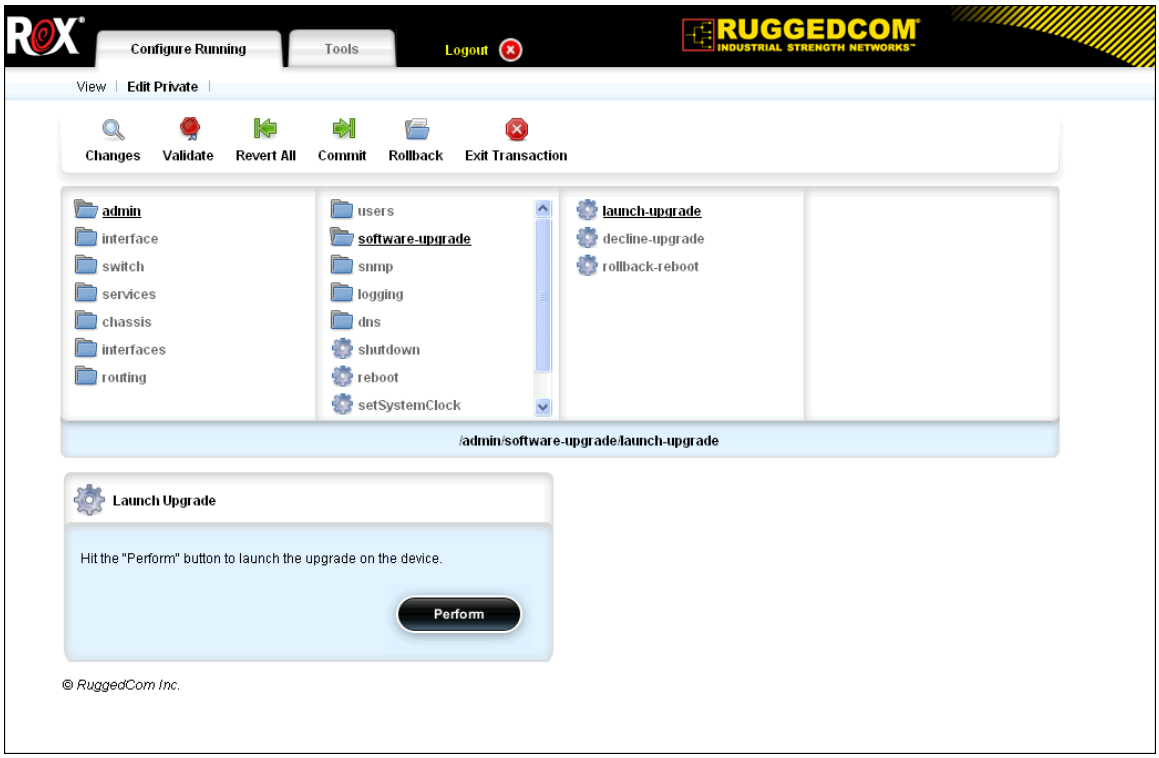

Figure A.5. Launch Upgrade

The Success! and Upgrade Options messages shown below indicate that the upgrade has been launched.

Success!
The upgrade has been performed.
Upgrade Options
Launch-result *
The upgrade has been launched. View 'upgrade-progress' to monitor its status.

Figure A.6. Upgrade Launched Dialogs

Click the **Exit Transaction** ( ) button at the top of the screen to return to the *View* mode.

A.3. Monitoring The Software Upgrade

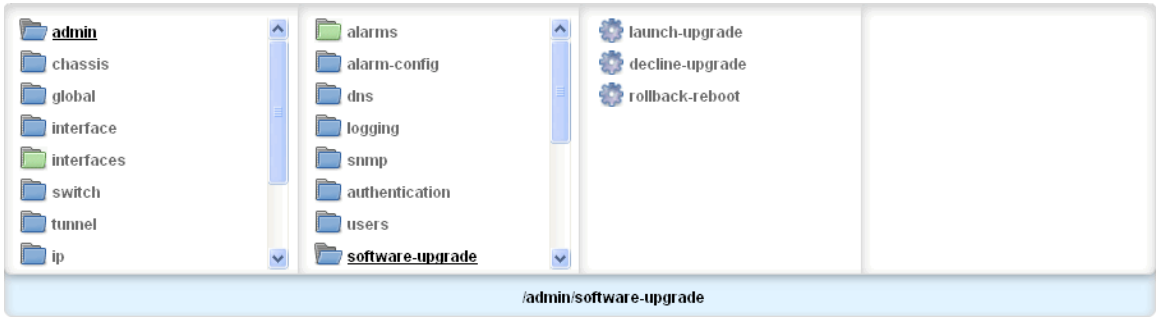

Figure A.7. Software-Upgrade Menu

Click on the Software-Upgrade menu to view the Upgrade Monitoring form. The Upgrade Monitoring form shows the real-time progress of the Upgrade procedure. The software upgrade progresses through four phases:

- Estimating upgrade size
- Copying filesystem
- Downloading packages
- Installing packages

The final reported status is either *Completed successfully* or, if an error prevented the completion of any one of the above four phases, *Failed*. These phases are shown in real time in the *Upgrade Phase* field on the Upgrade Monitoring Form, below.

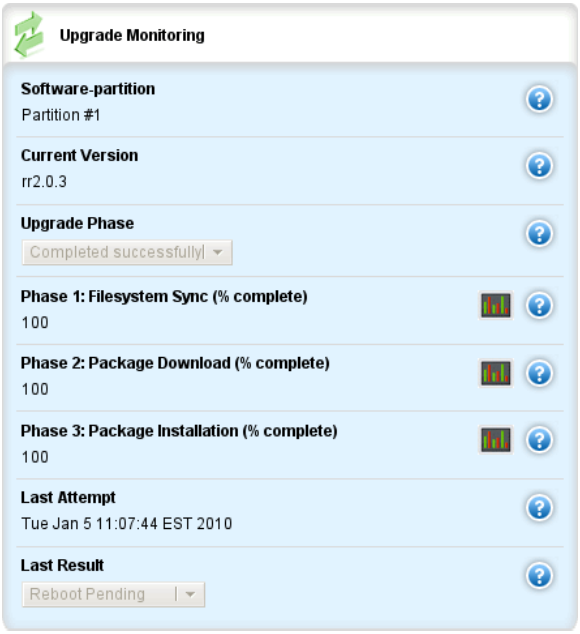

Figure A.8. Upgrade Monitoring Form in Reboot-pending Stage

Once the upgrade has successfully installed, the three phase progress indicator fields indicate 100% complete. In this state, the Upgrade Monitoring form will display *Reboot Pending* in the *Last Result* field. This indicates that a system reboot is required in order for the newly installed software upgrade to take effect.

Reboot the system.

When system has been rebooted to run the newly upgraded software installation, the *Current Version* field will reflect the ROX version number:

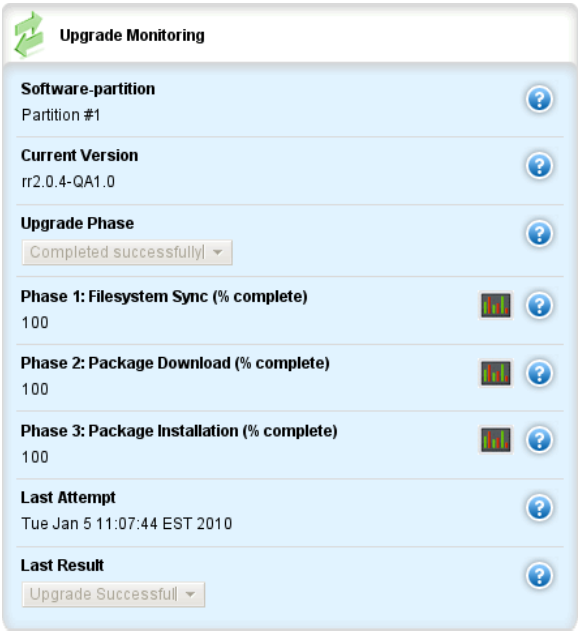

Figure A.9. Upgrade Monitoring Form Showing Successful Upgrade

software-partition

synopsis: a string of at most 31 characters

The current active partition number. The unit has two software partitions: #1 and #2. Upgrades are always performed to the other partition.

current-version

synopsis: a string of at most 31 characters

The current operating software version.

upgrade-state

synopsis: string - one of { Failed, Completed successfully, Unknown state, Installing packages, Downloading packages, Copying filesystem, Estimating upgrade size, Inactive }

The current phase or state of the upgrade.

filesystem-copy

synopsis: integer in the range [0 to 100]

Phase 1 of the upgrade involves synchronizing the filesystem with the partition to which we are upgrading.

This reflects the estimated percent complete.

package-download

synopsis: integer in the range [0 to 100]

Phase 2 of the upgrade downloads all packages that require an update. This reflects the estimated percent complete.

package-installation

synopsis: integer in the range [0 to 100]

Phase 3 of the upgrade installs all packages that require an update. This reflects the estimated percent complete.

last-upgrade-attempt

synopsis: a string of at most 64 characters

The date and time of completion of the last upgrade attempt.

last-upgrade-result

synopsis: string - one of { Interrupted, Declined, Not Applicable, Reboot Pending, Unknown, Upgrade Failed, Upgrade Successful }

Indicates whether or not the last upgrade completed successfully

Appendix B. RADIUS Server Configuration

This section describes the configuration procedures for two popular RADIUS servers, "FreeRADIUS" and the Microsoft Windows "Internet Authentication Service" in order to create and manage accounts that are able to access resources on RuggedBackbone™. There are four RADIUS attributes required for the configuration of accounts to access services on RuggedBackbone™. The following table shows the RADIUS attributes required by RuggedBackbone™ for accounts that are designated to use one or more of the "login", "ppp", or "ssh" services:

RADIUS Attribute	login	ppp	ssh
User ID	required	required	required
Password	required	required	required
NAS-Identifier			
RuggedCom-Privilege-level			

Table B.1. Required Attributes for various RADIUS services

Every account to be authenticated on behalf of the RuggedBackbone™ must have a user ID and password. The RADIUS "NAS-Identifier" attribute may optionally be used to restrict which service an account may access:

- *login*
- *ppp*
- *ssh*

Accounts that do not specify a "NAS-Identifier" attribute may access any RuggedBackbone™ service upon authentication. Accounts may also be defined to have access to one or several services. For more information on these services on RuggedBackbone™, please refer to RADIUS, ROXII, and Services.

You must all the following information to the vendor-specific extensions of the chosen RADIUS server:

- RuggedCom uses Vendor number 15004.
- "RuggedCom-Privilege-level" is attribute 2, of type "string".
- "RuggedCom-Privilege-level" must take one of the following three values:
  - "admin"
  - "operator"
  - "guest"

B.1. PPP / CHAP and Windows IAS

In order for Windows IAS to authenticate PPP connections that use the CHAP authentication protocol, IAS must be made to store passwords using what it calls "reversible encryption".

1. Ensure that CHAP authentication is enabled in the Remote Access Policy.
2. In the Active Directory settings for each PPP user, select "Store password using reversible encryption".

Appendix C. Setting Up An Upgrade Server

The RuggedCom software upgrade mechanism requires a repository of software to be available. The following instructions detail:

- Requirements for a repository server,
- Initial set up of a repository,
- Upgrading the repository to the latest release,
- Maintain separate releases streams for different groups of routers,
- Setting up one router to test new releases
- Configuring the network routers.

C.1. Upgrade Server Requirements

In order to establish a repository you will need a host that is accessible to the units that will be upgraded. This host must be able to act as a web server or ftp server. The host must also be able to access the RuggedCom web site in order to download new releases of software from RuggedCom.

The server requirements are fairly modest. The principal requirements are for disk space, bandwidth and the ability to serve an adequate number of http sessions.

Each software release will require approximately 75 Mb of disk space. Note that this figure includes an entire software image, most upgrades will involve the transfer of only a small fraction of this amount. A large number of such releases could easily be stored on a system of only modest capabilities. In practice, only one or two releases are usually all that need be kept.

The bandwidth requirements are determined by the many factors including the number of units, size of upgrade, when the routers upgrade, each unit's upgrade is bandwidth limited to 500kbps by default. Most web servers can serve files to the limit of the network interface bandwidth, so even a modest (e.g. 486 class machine) would prove acceptable.

The server should be able to accept at least as many http or ftp connections as there are upgradable units in the network. In practice you will configure the units to have staggered upgrade times in order to minimize the impact of upgrading on the network. A large upgrade (or a low bandwidth limiting value at each router) may cause all the routers to be upgrading at any one time.

C.2. Initial Upgrade Server Setup

You must create a directory on the web server to hold the releases for the router. The directory can have any name, such as “ruggedBackbone”.

Some administrators like to test the upgrade to the new release before making it available at the repository-url to which all the routers on their network are pointing. Perhaps this is desired if you have automated the routers to run the upgrade periodically, as launching an upgrade to a repository with the same release returns with no action. In this case simply create a separate directory in the webserver such as "ruggedBackboneTest", so the new release is not available to the routers on your network.

These directory names will be used in examples in the remainder of this section.

Ensure that the web server publishes these directories.

C.3. Upgrading The Repository

Releases are obtained from the RuggedCom web site as ZIP files. Download the ZIP file to your regular and/or test release directories and unzip them. You may delete the original ZIP file if desired.

The ZIP file name will be in the form rrX.Y.Z.zip. The major release number ‘X’ is changed when major new functionality (often hardware related) is offered. The minor release number ‘Y’ is increased when new features are added or serious bugs fixed, and the patch release number ‘Z’ is increased when minor bug repairs are made.

The zip file will extract to a directory that has the same name as the major release, e.g. “rr2”. As subsequent releases are made, they will also be extracted into this directory.

C.4. Setting Up The Routers

Suppose you have just unzipped rr2.1.1.zip into “ruggedBackbone” (or "ruggedBackboneTest" if you do not want your routers to point to it yet) on a server available to the network at server.xyz.net.

There are now two approaches available to you for the upgrade settings on the unit.

The first method described allows you to configure the unit to always upgrade to the last release that was installed on the server. The advantage of this approach is that you do not have to update upgrade configurations each time you obtain a new release. On the RuggedBackbone configure the *repository-url* parameter to “http://server.xyz.net/ruggedBackbone/rr2” and the *target-release* parameter to “current”. You can proceed to launch the upgrade on the ruggedBackbone (this can be done via CLI, web user interface, and NETCONF - see the appropriate user guide for details).

The second method allows you to configure the target release version explicitly. Some administrators may prefer this approach for sake of clarity, but this has the disadvantage of having to update all the units with the release version each upgrade. On the RuggedBackbone configure the *repository-url* parameter to “http://server.xyz.net/ruggedBackbone/rr2” and the *target-release* parameter to "rr2.1.1".

C.5. Using Microsoft Internet Information Services (IIS) Manager 6.0 or Higher as a ROX Upgrade Repository

When using Microsoft Internet Information Services (IIS) Manager 6.0 or higher as your ROX upgrade repository, you must add a new application/octet-stream MIME type named ***** to the IIS properties. The new MIME type is required for IIS to consider ROX upgrade packets as an application/octet stream. If the new MIME type is not added, ROX upgrades will fail.

To add the new MIME type, perform the following steps on your IIS server.

Procedure C.1. Adding the * MIME Type

1. In the Windows **Start** menu, right-click on **My Computer** and select **Manage**. The *Computer Management* dialog appears.
2. Under **Services and Applications**, locate the **Internet Information Services (IIS) Manager** node. Right-click on your ROX upgrade repository website and select **Properties**. The *Properties* dialog appears.
3. Select the **HTTP Headers** tab and click **MIME Types**. The *MIME Types* dialog appears.
4. Click **New**. The *MIME Type* dialog appears.
5. In the *Extension* field, type *****.

   In the *MIME type* field, type **application/octet-stream**.
6. Click **OK** on the *MIME Type*, *MIME Types*, and *Properties* dialog boxes.

Appendix D. Adding and Replacing Line Modules

Procedures for Adding and Replacing Line Modules

ROX™ version 2.2 does not support full hot-swap capability of line modules. Please adhere to the following procedures when adding or replacing line modules.

D.1. Shutting Down the Unit to Remove/Insert Modules

An action called 'shutdown' is available under *admin* to shut down the RuggedBackbone™. After invoking the shutdown action, you may ascertain that the operating system has been halted by the following message on the serial console: "System Halted, OK to turn off power" or by the status LEDs of all line-modules turning off (the alarm LED remains on). You may remove/insert modules in this state (powered on but halted). You have at least 60 seconds before the unit will automatically boot. If this is insufficient time for you to perform insertions/removals, please remove power to the unit. If removing wiring is inconvenient, you may turn off power by removing the power-module(s) themselves.

D.2. Adding a Module to an Empty Slot

1. Ensure that ‘module-type’ for the empty slot (under *chassis/line-modules/*) is set to “none”; this allows the system to auto-detect the new module on next boot.
2. Shut down the RuggedBackbone™.
3. Insert the new module into the slot and boot the unit.
4. After boot-up, the new line-module is auto-detected and operational.
5. Under *interface*, interfaces have now been created for the new module; you may proceed with additional configurations for the module.

If step 1 is omitted, you may manually configure the 'module-type' of the new module on boot-up. After committing the module-type configuration, the line-module will power on, but its status LED will be red, indicating that it is not passing traffic. On the next boot, the module will be fully integrated and operational.

D.3. Swapping a Module with an Identical Backup Module

You may wish to replace a module with an identical backup module in the event it is damaged or a defect is discovered. This procedure assumes the slot is already configured for the module.

1. Shut down the RuggedBackbone™.
2. Replace module with the new module and boot the unit.
3. The new module is operational after boot-up.

D.4. Swapping a Module with a Different Type of Module

1. Shut down the RuggedBackbone™.
2. Replace the module and boot the unit.
3. After the unit boots, log in, and (under *chassis/line-modules*) configure 'module-type' to the new module and admin-enable it. Please note that changing the module-type will result in losing interface configurations for the line-module.
4. Commit your modification. In some cases, you may encounter alarms reporting "Internal Configuration Error" after this commit, they can be safely ignored and cleared in this context; they are artifacts of unsupported hot-swap capability in ROX™.

5. After the commit, the module will power on, but its LED will be red indicating it is not yet passing traffic as it is not fully integrated into the system.
6. Reboot the unit.
7. After boot-up, the module will be integrated and operational. Under *interface*, interfaces have now been created for the module; you may proceed with related configurations.

D.5. Swapping a Module with a Different Type of Module

1. Set the 'module-type' (under *chassis/line-modules* to "none"; this allows the system to auto-detect the new module on next boot. Note that after committing this modification, the module will power down.
2. Shut down the RuggedBackbone™.
3. Replace the old module with the new module and boot the unit.
4. After boot-up, the new module is auto-detected and operational.
5. Under /interface/, interfaces have now been created for the new module; you may proceed with additional configurations.

Appendix E. GNU General Public License

Version 2, June 1991
Copyright © 1989, 1991 Free Software Foundation, Inc.

Free Software Foundation, Inc.
51 Franklin Street, Fifth Floor,
Boston, MA 02110-1301
USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.
Version 2, June 1991

E.1. Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software - to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps:

1. copyright the software, and
2. offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

E.2. TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

E.2.1. Section 0

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The “Program”, below, refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification”.) Each licensee is addressed as “you”.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

E.2.2. Section 1

You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

E.2.3. Section 2

You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: If the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

E.2.4. Section 3

You may copy and distribute the Program (or a work based on it, under Section 2 in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

E.2.5. Section 4

You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

E.2.6. Section 5

You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

E.2.7. Section 6

Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these

terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

E.2.8. Section 7

If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

E.2.9. Section 8

If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

E.2.10. Section 9

The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

E.2.11. Section 10

If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

E.2.12. NO WARRANTY Section 11

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

E.2.13. Section 12

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

E.3. How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.> Copyright (C) <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type “show w”. This is free software, and you are welcome to redistribute it under certain conditions; type “show c” for details.

The hypothetical commands “show w” and “show c” should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than “show w” and “show c”; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program “Gnomovision” (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989 Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.